ISSN 2039-6651
anno 2026 numero 113

ML4SEG

MISCELLANEA INGV

Abstract Volume

ML4SEG 2026 - Machine Learning for Solid Earth Geosciences and Earthquake Physics

Baia Samuele, Sicily | 14 - 19 June 2026

ISTITUTO NAZIONALE DI GEOFISICA E VULCANOLOGIA

113



Direttore Responsabile
Daniela VERSACE

Editor in Chief
Milena MORETTI

Editorial Board
Laura ALFONSI
Sofia BARANELLO
Carla BARNABA
Christian BIGNAMI
Simona CARANNANTE
Maria Grazia CIACCIO
Luca COCCHI
Luigi CUCCI
Lorenzo CUGLIARI
Elisabetta D'ANASTASIO
Alessia DI CAPRIO
Roberto DI MARTINO
Domenico DI MAURO
Domenico DORONZO
Filippo GRECO
Alessandro IAROCCI
Mario MATTIA
Daniele MELINI
Anna NARDI
Lucia NARDONE
Marco OLIVIERI
Arianna PESCI
Stefano PINTORE
Pierangelo ROMANO
Paola SBARRA
Davide SCAFIDI
Maurizio SOLDANI
Sara STOPPONI
Umberto TAMMARO
David ZULIANI

REGISTRAZIONE AL TRIBUNALE DI ROMA N.178 | 2014, 23 LUGLIO
© INGV Istituto Nazionale di Geofisica e Vulcanologia
Rappresentante legale: Presidente INGV
Sede: Via di Vigna Murata, 605 | Roma
Tutti i diritti riservati

www.ingv.it

ISTITUTO NAZIONALE **DI GEOFISICA E VULCANOLOGIA**

ISSN 2039-6651
anno **2026** numero **113**

MISCELLANEA INGV

Abstract Volume

ML4SEG 2026 - Machine Learning for Solid Earth Geosciences and Earthquake Physics

Baia Samuele, Sicily, Italy | 14 - 19 June 2026

Editors: Flavio Cannavò[1,*], Paul Johnson[2], Chris James Marone[3], Christopher William Johnson[4], Maurizio Petrelli[5], Andrea Cannata[6], Concetto Spampinato[7]

[1]INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Catania, Italy
[2]Los Alamos National Laboratory, Earth and Environmental Sciences Division, Los Alamos, NM, USA
[3]Sapienza University of Rome, Department of Earth Sciences, Rome, Italy
[4]Los Alamos National Laboratory, Earth and Environmental Sciences Division: Geophysics, Los Alamos, NM, USA
[5]University of Perugia, Department of Physics and Geology, Perugia, Italy
[6]University of Catania, Department of Biological, Geological, and Environmental Sciences, Catania, Italy
[7]University of Catania, Department of Electrical, Electronic, and Computer Engineering, Catania, Italy

**Corresponding author*

Accepted 6 August 2026 | *Accettato 6 agosto 2026*

How to cite | *Come citare* F., Cannavò, P., Johnson, C., James Marone, C., William Johnson, M., Petrelli, A., Cannata, C., Spampinato (Eds.) (2026). ML4SEG 2026 - Machine Learning for Solid Earth Geosciences and Earthquake Physics. *Miscellanea INGV*, *113*, 78 pp. https://doi.org/10.13127/misc/113

Cover Graphyc by B. Angioni (from AdobeStock 522724433 standard license) | *In copertina Elaborazione grafica di B. Angioni (da AdobeStock 522724433 licenza standard)*

113

Organising Committee

Flavio Cannavò
Chris James Marone
Paul Johnson
Christopher William Johnson
Maurizio Petrelli
Andrea Cannata
Concetto Spampinato

Local Logistics

Barbara Iurato (Eracle)
Veronica Cerruti (Eracle)

Invited Keynote Speakers

Prof. Gregory Beroza
Prof. Ivan Dokmanić
Prof. Yehuda Ben-Zion
Dr. Velimir Vessilinov
Dr. Daniel O'Malley
Dr. Laura Laurenti
Prof. Jonathan Bedford
Prof. Shokouhi Parisa
Prof. Maarten V. de Hoop
Dr. Adam Goldstein

Sponsors & Partners

CENTRO DI RICERCA
PREVISIONE PREVENZIONE E CONTROLLO
DEI RISCHI GEOLOGICI - CERI

SAPIENZA
UNIVERSITÀ DI ROMA

erc
European Research Council
Established by the European Commission

MUR
Ministero dell'Università e della Ricerca

FONDO ITALIANO PER LA SCIENZA

ISTITUTO NAZIONALE
DI GEOFISICA E VULCANOLOGIA

SICILIAE STVDIVM GENERALE 1434
Università di Catania

STUDIUM GENERALE CIVITATIS PERUSII A.D. MCCCVIII
A.D. 1308
unipg
DEPARTMENT
OF PHYSICS AND GEOLOGY

INDEX

Preface 11
The Organizing Committee

Beyond counting events: what makes a good earthquake catalog? 13
Gregory C. Beroza, Albert Leonardo Aguilar Suarez, Jaehong Chung, Yifan Yu, Lauro Chiaraluce, Maddalena Michele

Detection of (a)seismic transients using unsupervised Machine Learning and Deep Learning 14
Patricia Martínez-Garzón, Zahra Zali, Sadegh Karimpouli, Gregory C. Beroza

Localization of low magnitude seismicity before larger earthquakes 15
Yehuda Ben-Zion

Stochastic in-context foundation models 16
Maarten V. de Hoop, Trevor Teolis, Takashi Furuya, Ivan Dokmanić

Warpology: mechanistic surrogates for waves, grains, and vortices 17
Till Muser, Alexandra Spitzer, Maarten de Hoop, Matti Lassas, Ivan Dokmanić

From fundamental earthquake rupture physics to Machine Learning infused ground-motion simulations 18
Paul Martin Mai

Toward generalizable seismic intelligence: from task-specific models to foundation models and explainable AI 19
Laura Laurenti, Fabio Galasso, Christopher Johnson, Brian Kulis, David Liu, Michele Magrini, Chris Marone, Francesco Marrocco, Men-Andrin Meier, Gabriele Paoletti, Elisa Tinti, Daniele Trappolini, Stefan Wiemer

Leveraging machine learning to illuminate magmatic and tectonic processes during the 2024 unrest at Kīlauea volcano 20
Federica Lanza, Alicia Rohnacher, Ninfa Bennington, Stefan Wiemer

From foreshock-aftershock classification to time-dependent seismic cycle monitoring: station-wise CNN-SHAP analysis of Norcia 2016 21
Michele Magrini, Francesco Marrocco, Laura Laurenti, Gabriele Paoletti, Elisa Tinti, Chris Marone

AI framework for ground-motion prediction across the Italian Seismic Network 24
Aurora Bassani, Daniele Trappolini, Alessandro Scifoni, Giulio Poggiali, Elisa Tinti, Fabio Galasso, Alberto Michelini, Chris Marone

Deep Learning and ray path analysis illuminate spatiotemporal variations in fault zone properties during the seismic cycle 25
Gabriele Paoletti, Daniele Trappolini, Elisa Tinti, Fabio Galasso, Cristiano Collettini, Chris Marone

Unsupervised encoding of seismo-volcanic signals of a full eruptive cycle at Piton de la Fournaise (2014-2023) 26
Marie A. Gärtner, Reza Esfahani, Michel Campillo, Nikolai Shapiro

Automated AI agent for earthquake location and magnitude estimation: application to Campi Flegrei caldera 2014-2025 seismicity 27
Andrea Sollai, Aldo Zollo, Grazia De Landro, Salvatore Cuomo

Quality evaluation of assimilation-based forecast of rate-and-state governed fault analog 28
Bharath Shanmugasundaram, Harsha Bhat, Romain Jolivet

Foundation models reveal seismic damage evolution in fault zones 29
Sadegh Karimpouli, Patricia Martínez-Garzón, Frederik Tilmann, Georg Dresen, Marco Bohnhoff

Machine Learning based prediction of laboratory fault peak slip-rate trends using seismicity catalog 30
Xu Han, Sadegh Karimpouli, Marco Bohnhoff, Georg Dresen

Masked graph neural network for rapid ground motion prediction in Italy 31
Ilaria Oliveti, Daniele Trappolini, Licia Faenza, Dario Jozinović, Alberto Michelini

AI-enabled digital twins for volcano monitoring: insights from the DT-GEO project 32
Flavio Cannavò, Paolo Papale, Deepak Garg, Martina Allegra, Rebecca Bruni, Erica De Paolo, Chiara Montagna, Gilda Currenti, Danilo Reitano, Michele Carpené, Piero Lanucara

Application of data science and Machine Learning techniques for the prediction of induced seismicity 33
Leticia Garay Romero, Licia Faenza, Alexander Garcia, Anna Maria Lombardi

Deep Learning-based urban acoustic source identification from DAS data in the time-space-frequency domain 34
Yichen Zhong, Chen Gu, Michael Fehler, German A. Prieto, Peng Wu, Zhi Yuan, Jiawei He, Chenrui Fan, Yufeng He

Explainable physics-guided feature extraction from time-lapse ultrasonic waveforms for predicting laboratory earthquakes 35
Prabhav Borate, Jacques Rivière, Daniel Kifer, Parisa Shokouhi

Beyond prospectivity maps: multimodal geospatial AI for critical mineral systems and geothermal prospectivity exploration 36
Velimir Vesselinov

Time-reversal neural operator for earthquake rupture inversion via data fusion of broadband seismic arrays and submarine DAS 37
Chen Gu, Zhuoyu Chen, Jiawei He, Yufeng He, Michael Fehler, Yichen Zhong, Peng Wu, Zhi Yuan, Chenrui Fan, Shigui Zhang, Jiachao Xu, Xiaoyong Qin, Xiaoyan Fan, Hongbo Sun

Mitigating atmospheric noise in InSAR time series over volcanoes using a convolutional neural network 38
Rebecca Bussard, Christelle Wauthier

In-context learning enables global geothermal temperature prediction from sparse local observations 39
Daniel O'Malley, David Coblentz, Nathan Debardeleben, Chris Johnson, John Kath, Pablo Lara, Earl Lawrence, Sandro Malusa, Arnab Mazumder, Mohamed Mehana, Javier E. Santos, Bharat Srikishan, Hari Viswanathan

The Transparent Earth: a multimodal foundation model for Earth sciences 40
Sandro Malusà, Javier E.Santos, Pablo Eduardo Espinoza Lara, Christopher Johnson, David Coblentz, Daniel O'Malley

Scale-aware seismic phase arrival detection via continuous dilatable convolutional kernels 41
Sacha Lapins

Towards DL phase picking for diverse volcanic seismicity: a case study at Tungurahua 42
Abigail Robinson, Prof. James Hammond, Dr Sacha Lapins, Dr Alex Steele

High-resolution lithospheric thickness model of the Southern Indian Shield using Physics-Informed Neural Networks and geophysical data integration 43
Arka Roy, Jörg Ebbing

Advanced data model for the multidisciplinary cataloguing of volcanic ash 44
Simona Caruso, Flavio Cannavò, Martina Allegra, Lucia Cacciola, Daniele Andronico, Antonio Cristaldi

AI-driven automated and scalable processing of DAS data for seismic monitoring 45
Miriana Corsaro, Flavio Cannavò, Gilda Currenti, Simone Palazzo, Martina Allegra, Philippe Jousset, Michele Prestifilippo, Concetto Spampinato

Deciphering pre-eruptive dynamics in Antarctic and Italian volcanoes through Machine Learning 46
Mónica Ágreda-López, Irene Rocchi, Pier Paolo Giacomoni, Maurizio Petrelli, Matteo Masotta, Sergio Rocchi

Towards Machine Learning-based characterization of seismo-volcanic events in a closed-conduit volcano: the case of Mt. Melbourne, Antarctica 47
Graziano Larocca, Danilo Contrafatto, Gaetano Giudice, Francesco Ciancitto, Daniele Pellegrino, Alessandro Bonforte, Salvatore Consoli, Francesco Panzera, Mariangela Sciotto, Andrea Cannata, Flavio Cannavò, Loredana Contrafatto, Simona Scollo

Early detection of paroxysmal eruptions at Stromboli volcano using seismic data 49
Vittorio Minio, Flavio Cannavò, Marcello D'Agostino, Massimo Orazi, Eugenio Privitera

Clustering of tiltmeter data for automated identification of Mount Etna's active crater 50
Adriana Iozzia, Marco Aloisi, Flavio Cannavò, Salvatore Gambino

Seismology on rivers: a Machine Learning approach to monitor extreme floods generated by cyclonic storms 51
Mario Valerio Gangemi, Alfio Marco Borzì, Martina Allegra, Andrea Cannata, Flavio Cannavò, Concetto Spampinato, Luca Cavallaro, Enrico Foti, Rosaria Ester Musumeci, Stefano Parolai, Luca Zini, Matteo Picozzi, Salvatore Alparone, Danilo Contrafatto, Graziano Larocca, Francesco Panzera

Reconstructing significant wave height from microseism data using Machine Learning techniques: Sicily Channel case study 53
Alfio Marco Borzì, Vittorio Minio, Andrea Cannata, Francesco Panzera, Susanna Saitta, Salvatore Alparone, Giuseppe Ciraolo, Danilo Contrafatto, Sebastiano D'Amico, Giuseppe Di Grazia, Graziano Larocca, Flavio Cannavò

Hidden simplicity behind seismic dispersion curves revealed by unsupervised learning 54
Han Zhang, Jiao Yang, Sidao Ni, Zefeng Li, Xuewei Bao, Lijun Liu, Brandon Schmandt

Physics-guided Machine Learning for meteotsunami detection from co-located microbarometer and tide-gauge records at Mazara del Vallo, Italy 55
Salvatore D'Amico, Salvatore Aronica, Andrea Cannata, Francesco Panzera

DeepPinch: few-shot stratigraphic pinch-out detection in seismic images via synthetic pre-training 57
Shimon Cohen, Amotz Agnon, Michael Werman

Crustal fracture classification and parameter estimation from InSAR displacement fields with interpretable descriptors 58
Louise Lizé, Christophe Lin-Kwong-Chon, Valérie Cayol

Using neural networks for InSAR images inverse modeling 59
Gregoire Vest, Yajing Yan, Fabien Albino

Template matching using Deep Learning 60
Dario Jozinović, Tania Toledo, Verena Simon, Toni Kraft, Frederick Massin, John Clinton, Amilcar Roca, Olivier Meyers

Unsupervised discrimination of earthquakes and blasts in Norway using representation learning 61
Miguel Neves, Lars Ottemöller, Stéphane Rondenay

Integrating game engine and PINN for multi-scale building cluster earthquake scenarios 62
Chenrui Fan, Chen Gu, Zhuoyu Chen, Yuheng Wu, Zhihua Zhao, Yichen Zhong, Jiawei He, Peng Wu, Zhi Yuan, Yufeng He, Yifu Cheng

Diffusion-based denoising for submarine DAS seismic monitoring 63
Jiawei He, Chen Gu, Yichen Zhong, Peng Wu, Zhi Yuan, Chenrui Fan, Shigui Zhang, Jiachao Xu, Xiaoyong Qin, Xiaoyan Fan, Hongbo Sun

From faults to volumes: the development of an invertible Deep Learning surrogate for versatile kinematic source modelling 64
Jonathan Bedford, Kaan Çökerim

Blind surrogate inversion of deformation sources 65
Kaan Çökerim, Jonathan Bedford

Source localization from submarine DAS data using deep Bayesian beamforming 66
HE Yufeng, GU Chen, HE Jiawei, CHEN Zhuoyu, ZHONG Yichen, WU Peng, YUAN Zhi, FAN Chenrui, WU Yuheng

Earthquake seismogram denoising across time, time–frequency, and hybrid domain approaches 67
Niko Dahmen

Time delay neural networks reveal pressure-independent fault rupture processes in laboratory acoustic emission 68
Sergio Vinciguerra, Thomas King

N-dimensional field reconstruction from sparse observations 69
Javier E. Santos

Autonomous research agents and structured feature operators for earthquake forecasting 70
Allie Hutchison, Ryan Parks, Daniel Friedman, Adam Goldstein

Adaptive Deep Learning approach for laboratory earthquake forecasting using multi-channel acoustic emission data 71
Ke Gao, Zhixiong Leng

GPU accelerated SEAS simulations for Machine Learning applications 72
Gabriele Benedetti, Elías Rafn Heimisson

Meta-learning for few-shot forecasting of laboratory earthquakes 73
Zhixiong Leng, Ke Gao

Automatic analysis of seismic waveforms for sea-ice properties estimation 74
Jonas Michael, Ludovic Moreau, Marielle Malfante

Preface

Welcome to the Book of Abstracts of ML4SEG 2026 – Machine Learning for Solid Earth Geosciences and Earthquake Physics, held in Baia Samuele, Sicily, Italy, from 14 to 19 June 2026.

ML4SEG was conceived as an international forum where geoscientists, physicists, engineers, computer scientists, and data scientists can meet to explore how modern machine learning is reshaping the study of the solid Earth. The conference aims to bridge recent advances in artificial intelligence with some of the most challenging questions in earthquake physics, volcano monitoring, geodynamics, geochemistry, geodesy, and geophysical exploration.

The contributions collected in this volume reflect the breadth and vitality of this rapidly evolving field. They range from deep-learning methods for seismic detection, phase picking, denoising, earthquake location, and ground-motion prediction, to physics-informed neural networks, foundation models, neural operators, digital twins, explainable AI, and data-assimilation approaches. Several abstracts also address volcanic unrest, eruption forecasting, laboratory earthquake experiments, distributed acoustic sensing, InSAR inversion, geothermal and mineral exploration, sea-state and flood monitoring, and the integration of physical models with data-driven approaches. Together, these works show that machine learning is no longer only a tool for automation, but is becoming a framework for scientific discovery, hypothesis testing, uncertainty quantification, and operational decision support in the geosciences.

A recurring theme throughout the conference is the need to combine predictive performance with physical understanding. Solid Earth systems are complex, nonlinear, multi-scale, and only partially observed. Their study requires methods that can learn from large datasets while remaining interpretable, robust, transferable, and consistent with physical principles. The abstracts presented here demonstrate how the community is addressing these challenges through hybrid approaches that connect observations, simulations, laboratory experiments, and theory.

ML4SEG 2026 also emphasizes collaboration across disciplinary boundaries. Progress in this field depends on open dialogue between domain experts who understand the physics of natural systems and machine-learning researchers who develop new algorithms, architectures, and computational strategies. By bringing these communities together, the conference seeks to encourage reproducible workflows, shared benchmarks, responsible use of AI, and practical tools that can contribute to monitoring, forecasting, hazard assessment, and a deeper understanding of Earth processes.

We thank all authors for their contributions, the invited speakers for sharing their perspectives, the scientific and organizing committees for their work, and all participants for helping to build a dynamic and inclusive scientific environment. We hope that this volume will serve not only as a record of the conference, but also as a source of ideas, collaborations, and future developments at the intersection of machine learning and solid Earth science.

The Organizing Committee

Keywords Machine Learning; Solid Earth; Geosciences

Beyond counting events: what makes a good earthquake catalog?

Gregory C. Beroza[1,*], Albert Leonardo Aguilar Suarez[1], Jaehong Chung[1], Yifan Yu[1], Lauro Chiaraluce[2], Maddalena Michele[2]

[1]*Stanford University, Stanford, CA, USA*
[2]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Osservatorio Nazionale Terremoti, Rome, Italy*

**Corresponding author: beroza@stanford.edu*

Recent advances in deep learning (DL) methods have created a growing number of earthquake catalogs that contain many more earthquakes than those assembled through conventional approaches. Label errors, hallucinations, and data leakage have the potential to undermine the reliability of DL results, and there is a growing need to understand the strengths and weaknesses of DL-based earthquake monitoring. Understanding potential shortcomings of current models is a necessary step in using them for scientific inference, in making them operational, and for improving them to realize additional performance gains. We need to move beyond counting events as a measure of effectiveness.

We lack absolute ground truth, but existing earthquake catalogs and the phases associated to construct them, are carefully quality controlled and provide an approximation. We can explore the effectiveness of DL by comparing the detection probabilities of previously associated phases with those that are picked through deep learning. For an earthquake catalog constructed for the 2016-2017 Amatrice-Visso-Norcia sequence, we find that the PhaseNet picker increases detection sensitivity by identifying smaller magnitude events and detecting earthquakes at greater distances; however, the DL phase picking also shows much greater performance variability across stations. For event detection, this translates into a probabilistic magnitude-of-completeness improvement from 1.6 to 0.5 for P waves and from 1.7 to 0.5 for S waves. Across nearly all published DL catalogs, the fraction of previously known events recovered by DL-based monitoring, known as the recall, is well under 1 and includes earthquakes across the size spectrum. These results demonstrate that DL-based improvements in detectability are widespread but non-uniform, highlighting their benefits, their limitations, and the potential for further improvement.

Detection of (a)seismic transients using unsupervised Machine Learning and Deep Learning

Patricia Martínez-Garzón[1,2,*], Zahra Zali[1], Sadegh Karimpouli[1], Gregory C. Beroza[3]

[1]*GFZ Helmholtz Centre for Geosciences, Potsdam, Germany*
[2]*RWTH Aachen University, Aachen, Germany*
[3]*Stanford University, Stanford, CA, USA*

**Corresponding author: patricia@gfz.de*

Intermediate-term transient fault behaviors—ranging from seismic to aseismic slip—have been postulated as potential precursors to large earthquakes. However, their detection remains challenging due to variability across sequences, and limited and heterogeneous monitoring network. In this work, we develop unsupervised and deep learning approaches to identify both aseismic and seismic transients embedded in continuous strainmeter and seismic observations. In the first part, we present a fully automated deep learning–based workflow to detect short-duration slow slip events (SSEs), ranging from minutes to hours, in continuous borehole strainmeter records. Borehole strainmeters provide highly sensitive measurements of crustal strain and therefore capture such subtle signals. Nevertheless, SSE detection is challenging because their small amplitudes can be easily masked by non-tectonic sources. To tackle this problem, we propose a framework with four steps: (1) preprocessing of raw strainmeter data to remove long-term trends, tidal signals, and barometric pressure effects; (2) applying a Morlet wavelet transform to enhance transient strain features; (3) extracting informative features using a deep convolutional autoencoder; and (4) categorizing of these features to identify SSEs as a distinct class separated from noise and other non-SSEs signals. We apply this workflow to eight years of data from three borehole strainmeters near Parkfield, California, at the edge of the creeping segment of the San Andreas Fault. The method identifies SSEs detected at all stations, with several events consistently recorded across multiple instruments and validated using a manual catalog and creepmeter observations. Preliminary locations suggest a shallow origin, and temporal analysis indicates a possible interaction with low-frequency earthquakes (LFEs) occurring on deeper portions of the fault. Their inferred duration and size suggest that these short-term SSEs follow a scaling law similar to that of earthquakes. In the second part, we introduce an unsupervised machine learning framework to categorize seismicity patterns and identify seismic transients that may precede large earthquakes. We analyze two events with well-documented preparatory phases: the 2009 M_W 6.1 L'Aquila earthquake (central Italy) and the 2023 M_W 7.8 Kahramanmaraş earthquake (southern Türkiye). In this framework, we extract features describing the seismo-mechanical properties assigned to event families—clusters of earthquakes grouped in space, time, and magnitude—and categorize them using unsupervised learning methods. The detected categories reflect distinct seismicity stages, ranging from background activity to a relatively critical phase preceding the mainshock. Results highlight the emergence of long-lasting families with distinctive properties during this critical stage, characterized by increased spatiotemporal localization, stronger event interactions, and enhanced strain release. In contrast, such a stage is not detected in cases where no preparatory seismic phase has been reported. The method is scalable from laboratory experiments to natural fault systems, and identifies characteristic patterns without prior assumptions, demonstrating potential for operational earthquake forecasting. Together, these complementary approaches illustrate how machine learning can uncover subtle transient seismic and aseismic fault processes. By systematically identifying these transients in continuous geophysical observations, such frameworks provide new opportunities to investigate fault dynamics and improve our understanding on how earthquakes start.

Localization of low magnitude seismicity before larger earthquakes

Yehuda Ben-Zion

University of Southern California, USA

Corresponding author: benzion@usc.edu

Large failure events in the crust and laboratory experiments involve localization of deformation from a volume with high strain energy onto a narrow rupture zone. Tracking the evolving localization can provide fundamental information on preparation processes and approaching times of large earthquakes. Analyses of background (declustered) seismicity revealed localization before several M>7 earthquakes in California, Turkey, and Alaska, followed by delocalization after the large events [Ben-Zion and Zaliapin, 2020; Zaliapin and Ben-Zion, 2023]. Using evolving Gini coefficients of the spatial distribution of background seismicity, these studies observed localization/delocalization cycles with maximum localization ~2-3 years before the M>7 earthquakes. However, the analysis was limited to the largest events in the catalogs and the final stages leading to the large earthquakes remained unobserved. Recent analysis using evolving median separation distances and times of failure events in laboratory experiments, induced seismicity, and tectonic earthquakes covering over 15 orders of magnitudes show that across all these scales seismicity progressively localizes toward eventual rupture zones and delocalizes afterward [Hou et al., 2026]. The observed spatiotemporal organization on the approach to relatively large failure events opens new possibilities for interpreting seismicity patterns and developing earthquake forecasting procedures.

References

Ben-Zion, Y., and I., Zaliapin (2019). Spatial variations of rock damage production by earthquakes in southern California, *Earth Planet. Sci. Lett.*, 512, 184–193, doi: 10.1016/j.epsl.2019.02.006

Ben-Zion, Y., and I., Zaliapin (2020). Localization and coalescence of seismicity before large earthquakes, *Geophys. J. Int.*, 223, 561–583, doi: 10.1093/gji/ggaa315

Hou, Y., H., Meng, Z., Liu, and Y., Ben-Zion (2026). Universal Space-Time Localization of Seismicity before Relatively Large Events across Laboratory, Industrial, and Tectonic Scales, in review.

Stochastic in-context foundation models

Maarten V. de Hoop*, Trevor Teolis, Takashi Furuya, Ivan Dokmanić

Rice University, USA

**Corresponding author: mdehoop@rice.edu*

Foundation models for physics are built around in-context learning: given observed states forming a context, the model must learn and evolve the relevant physical law. We propose a kinetic framework in which tokens are lifted to particles constituting an empirical measure. The dynamics are driven by externally sampled candidate jump histories, decoupled from the initial particle distributions through thinning. The in-context mechanism then enters through a jump acceptance gate and interaction rules. As a result, the same sampled history induces different particle trajectories, which define in-context maps for different initial distributions. The model output is obtained by taking the expectation over the jump histories of the induced measure evolutions. This construction connects transformer-like architectures with kinetic equations and stochastic jump processes.

At the particle level, the jumps can be modeled by a piecewise-deterministic Markov process (PDMP). At each accepted jump time, a collision is simulated between a tracked particle and a collision partner. This admits a Poisson-driven representation: Poisson random measures generate candidate events, while a state-dependent acceptance gate ensures that collisions occur at the correct rate. Through regularization, the induced in-context maps can be approximated using self-attention. The resulting architecture is not merely a transformer: its output is an expectation over stochastic histories, so the predicted measure can spread beyond the initial empirical support and represent an ensemble of possible interaction pathways.

As the regularization is removed, in the mean-field regime, the measure dynamics converge to solutions of Boltzmann-type equations: the random jumps produce, at the level of the limiting particle distribution, an effective binary collision operator. Thus, although the model is driven by sampled stochastic histories rather than by a deterministic collision map, its mean-field limit captures Boltzmann collision dynamics. This suggests a lift–evolve–project architecture for physics foundation models: lift physical data into a particle or measure representation, evolve it through the learned expectation of stochastic in-context maps, and project back to observable quantities. This foundation model is probably suited for learning physics such as granular flows and fault dynamics. This also suggests possible augmentations of the Nanbu dynamics with Hawkes-type memory mechanisms, in which past collision events influence future interaction rates.

Warpology: mechanistic surrogates for waves, grains, and vortices

Till Muser[1,*], Alexandra Spitzer[1], Maarten de Hoop[3], Matti Lassas[2], Ivan Dokmanić[1]

[1]*University of Basel, Department of Mathematics and Computer Science, Basel, Switzerland*
[2]*University of Helsinki, Department of Mathematics and Statistics, Helsinki, Finland*
[3]*Rice University, Department of Computational Applied Mathematics and Operations Research, Houston, TX, USA*

**Corresponding author: ivan.dokmanic@unibas.ch*

What do we want from neural surrogates when the goal is to do science? They should be fast, expressive, and interpretable, but also able to capture physically distinctive phenomena like localization, coherent transport, wave scattering, and vortex interactions. The question is how to represent physics so that these qualities come together. In this talk I will sketch how phase-space thinking, in the spirit of Hamilton and Boltzmann, can organize effective learned models for fluids, waves, and granular flows. I will then present Flowers, a family of neural surrogates built from nothing but learned coordinate changes, or warps, rather than attention, convolution, or Fourier multipliers. These unassuming warps perform strongly across diverse physics while remaining lightweight and directly visualizable. I will close with preliminary work on native spherical models for global circulation and some first experiments on granular dynamics.

From fundamental earthquake rupture physics to Machine Learning infused ground-motion simulations

Paul Martin Mai

King Abdullah University of Science and Technology, Earth Systems Science and Engineering Program Division of Physical Sciences and Engineering, Saudi Arabia

Corresponding author: martin.mai@kaust.edu.sa

Machine Learning and Artificial Intelligence (ML/AI) are nowadays ubiquitously applied in many fields of science and technology. In the field of earthquake seismology, ML-techniques were initially used for earthquake detection, earthquake location, estimation of potential future shaking levels, and other automatable applications, but only in a limited manner in the context of replicating the physics of earthquake rupture or generating the intricate high-frequency seismic wavefield affected by source complexity, 3D Earth structure and topography as well as local site effects.

In this presentation, I briefly recap the physics and related 3D simulations of earthquake rupture dynamics under spatially variable stress, friction, geometry, and material properties. Examining the results of a database of HPC-enabled dynamic-rupture simulations, I introduce a ML-enabled rupture generator that effectively reproduces the spatiotemporal source complexities seen in real earthquakes.

Pairing this rupture generator with an innovative FNO- and GAN-driven generative high-frequency seismic-wave simulator provides an avenue for fast, efficient, and reliable physics-based ground-motion simulations that circumvents the computational bottleneck of full 3D dynamic simulations up to frequencies of engineering interest. While still being under development and testing, one of the key challenges for any such approach is how to validate and then statistically rigorously integrate such simulations into common probabilistic seismic hazard assessment (PSHA).

Toward generalizable seismic intelligence: from task-specific models to foundation models and explainable AI

Laura Laurenti[1,*], Fabio Galasso[2], Christopher Johnson[3], Brian Kulis[4], David Liu[1], Michele Magrini[5], Chris Marone[6,7], Francesco Marrocco[5], Men-Andrin Meier[1], Gabriele Paoletti[6], Elisa Tinti[6], Daniele Trappolini[2], Stefan Wiemer[1]

[1]*Swiss Seismological Service (SED), ETH Zurich, Zurich, Switzerland*
[2]*Sapienza University of Rome, Department of Computer, Control and Management Engineering, Rome, Italy*
[3]*Los Alamos National Laboratory, Los Alamos, NM, USA*
[4]*Boston University, Boston, MA, USA*
[5]*Sapienza University of Rome, Department of Mathematics, Rome, Italy*
[6]*Sapienza University of Rome, Department of Earth Sciences, Rome, Italy*
[7]*Pennsylvania State University, Department of Geosciences, University Park, PA, USA*

**Corresponding author: laura.laurenti@sed.ethz.ch*

Machine Learning is rapidly transforming scientific research, and seismology is increasingly benefiting from this shift toward data-driven modeling. Traditional supervised approaches typically rely on task-specific models trained on labeled datasets, achieving strong performance on narrow objectives but often lacking generalization and robustness across diverse seismic settings.

More recently, foundation models have emerged as a promising alternative. By learning generalized representations from large-scale seismic datasets through self-supervised learning, these models enable a unified framework for multiple downstream tasks. Their learned embeddings are applied to tasks such as fault state classification, ground motion regression for early warning, and earthquake phase detection. Results demonstrate improved performance compared to task-specific models, even in few-shot settings, highlighting the promise of foundation models for seismology.

Despite this progress, seismic foundation models remain far from the scale, versatility, and general reasoning capabilities of foundation models seen in domains such as natural language processing. Significant challenges remain, including limited data diversity, uncertainty in model transferability, and the black-box nature of deep learning models. We will discuss these current limitations and highlight the role of explainable AI in improving model interpretability, trust, and scientific insight.

Together, these efforts contribute toward more robust, generalizable, and interpretable machine learning frameworks for seismic applications.

Leveraging machine learning to illuminate magmatic and tectonic processes during the 2024 unrest at Kīlauea volcano

Federica Lanza[1,*], Alicia Rohnacher[1], Ninfa Bennington[2], Stefan Wiemer[1]

[1]*Swiss Seismological Service (SED), ETH Zurich, Zurich, Switzerland*
[2]*U.S. Geological Survey, Hawaiian Volcano Observatory, Hawai'i, USA*

**Corresponding author: federica.lanza@sed.ethz.ch*

Kīlauea on Island of Hawai'i is one of the most active volcanoes worldwide and frequently experiences major volcanic eruptions and tectonic activity. Here, we apply a machine-learning-based detection and localization framework, *qseek*, to seismic data recorded during a three-month deployment of a dense nodal array of 116 stations located along the East Rift Zone of Kīlauea. The resulting high-resolution earthquake catalog achieves substantially improved completeness compared to routine monitoring catalogs, unlocking previously under-resolved regions of the lower East Rift Zone and revealing detailed spatiotemporal patterns of seismicity during the period of observation.

The ML-derived catalog captures the full sequence of the 2024 Nāpau volcanic crisis, including three dike intrusions, with the last intrusion culminating in an eruption on 16 September. The enhanced detection capability of this network allows us to track the spatial and temporal evolution of seismicity and infer the characteristics of magma transport within the ERZ during the pre- and co-eruptive phases. In addition, we identify earthquake clusters along the south flank following dike emplacement, highlighting the interaction between magmatic intrusions in the rift zone and tectonic activity along the decollement.

These observations underline the potential of dense nodal networks that, along with deep-learning catalogs, improve our understanding of magma migration and volcano-tectonic processes.

From foreshock-aftershock classification to time-dependent seismic cycle monitoring: station-wise CNN-SHAP analysis of Norcia 2016

Michele Magrini[1,*], Francesco Marrocco[1], Laura Laurenti[2], Gabriele Paoletti[3], Elisa Tinti[3], Chris Marone[3,4]

[1]*Sapienza University of Rome, Department of Mathematics, Rome, Italy*
[2]*Swiss Seismological Service (SED), ETH Zurich, Zurich, Switzerland*
[3]*Sapienza University of Rome, Department of Earth Sciences, Rome, Italy*
[4]*Pennsylvania State University, Department of Geosciences, University Park, PA, USA*

**Corresponding author: magrini.2066963@studenti.uniroma1.it*

Laboratory, theoretical, and field data suggest that fault-zone properties evolve during the seismic cycle as stress accumulates prior to failure and drops during earthquake rupture. Laboratory experiments show systematic changes in elastic properties during the seismic cycle and demonstrate that these changes can be used to predict laboratory earthquakes. Recent work indicates that, in some cases, similar approaches may also be applicable to tectonic faults. Seismic observations reveal clear differences between fault-zone properties before and after large earthquakes, as well as post-seismic time-dependent changes in elastic behavior. Here, we develop predictive and physically interpretable tools to investigate seismic waveforms recorded around the M6.5 2016 Norcia earthquake. We train convolutional neural networks, CNNs, on RGB spectrograms derived from three-component seismograms recorded at individual seismic stations surrounding the fault system. In previous work, we showed that CNN models can distinguish foreshocks from aftershocks and that SHapley Additive exPlanations, SHAP, can identify class-dependent time-frequency signatures in the spectrograms. We also found that models trained on multiple stations can learn discriminative features from spatially distributed seismic data, but that SHAP attribution patterns may be substantially reorganized with respect to station-wise models, highlighting the importance of station-dependent propagation effects. In this work, we focus on the station-wise setting in order to isolate local, station-specific information and obtain more stable physical interpretations. We first train station-wise CNN models for the binary classification of foreshocks and aftershocks. We then extend the same framework to a more detailed temporal classification problem, in which the seismic sequence is divided into eight chronological classes, each corresponding to a two-week interval. This allows us to move beyond a simple pre-/post-mainshock separation and to explore whether the learned spectrogram features contain information about time to/since failure and progressive evolution of the seismic sequence. For both the binary and the eight-class temporal classification tasks, we apply SHAP to identify the time-frequency regions that contribute most strongly to the CNN predictions. In the binary case, SHAP highlights class-dependent spectrogram signatures that separate foreshock and aftershock waveforms. In the eight-class setting, SHAP provides a time-dependent view of feature importance across successive stages of the sequence, revealing how the model shifts its attention among different spectrogram regions as the seismic cycle evolves. This interpretation framework offers a way to investigate whether the discriminative information learned by CNNs reflects progressive evolution of the fault zone state of stress prior to the mainshock, gradual post-seismic evolution, or station-specific observational effects.

Our results indicate that station-wise CNN models, combined with SHAP analysis, can provide an interpretable framework for studying seismic-cycle-dependent changes in waveform

properties. By extending the classification from a binary foreshock/aftershock task to an eight-class temporal problem, this work takes a step toward tracking the evolution of fault-zone properties through time rather than treating the mainshock as a single classification boundary. The findings motivate future work on physically constrained and station-aware interpretability methods for seismic monitoring.

AI framework for ground-motion prediction across the Italian seismic network

Aurora Bassani[1,*], Daniele Trappolini[1], Alessandro Scifoni[1], Giulio Poggiali[2], Elisa Tinti[2], Fabio Galasso[1], Alberto Michelini[3], Chris Marone[2,4]

[1]*Sapienza University of Rome, Department of Computer, Control and Management Engineering, Rome, Italy*
[2]*Sapienza University of Rome, Department of Earth Sciences, Rome, Italy*
[3]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Rome, Italy*
[4]*Pennsylvania State University, Department of Geosciences, University Park, PA, USA*

**Corresponding author: aurora.bassani@uniroma1.it*

Estimating ground-motion intensity at individual seismic stations is a fundamental task in seismology, with direct implications for seismic hazard assessment. Ground-motion intensity measures (IMs) such as Peak Ground Acceleration (PGA), Peak Ground Velocity (PGV), and Spectral Acceleration (SA) at selected periods are commonly used to quantify shaking severity and relate it to building structural response and seismic hazard.
Here, we present an AI-driven framework for predicting multiple IMs at the station level, building on previous graph-based approaches for the Italian seismic network.
Starting from the INSTANCE dataset, we applied filters on event magnitude (≥ 3) and waveform quality, resulting in 3076 events recorded across 565 stations. For waveform analysis, we used a 10-second time window starting 1 second before the first P-wave arrival, balancing prediction speed and accuracy.
Seismic waveforms are first encoded using a pre-trained PhaseNet model to extract compact temporal representations. Spatial dependencies are modeled with a masked Graph Convolutional Network (GCN) based on Delaunay triangulation, which links each station to its nearest neighbors while avoiding long or crossing edges. This structure allows identification of border stations (at mesh edges) and coastal stations (within 15 km of the coastline). A binary mask distinguishes these nodes, helping the model account for areas with high azimuthal gap, which can make IMs estimation more challenging. Before the final prediction layer, the model concatenates the maximum waveform amplitude across stations with event metadata predicted by a fine-tuned LLM for magnitude and location estimation. This enables joint exploitation of temporal waveform features, network geometry, and global event information.
The framework predicts PGA, PGV, and SA at periods of 0.3, 1, and 3 s at all stations. We obtained preliminary results for multiple configurations. The baseline model served as reference, which included waveform representations and the GCN but not LLM metadata, border features, or weighted loss. Including LLM-derived metadata consistently improved performance across all regions and parameters, reducing relative errors by 14% in Southern Italy and over 27% in Northern Italy. The addition of border and coastal features provided only minor improvements, confirming that metadata was the main factor driving gains.
Applying a weighted loss emphasizing stations closer to the epicenter further improved predictions, particularly in Northern and Southern Italy. In Central Italy, where network coverage is denser, improvements were smaller, suggesting that local contributions were already well captured. For the best configuration (LLM metadata and weighted loss), global mean absolute errors across stations were 0.511 (PGA), 0.423 (PGV), 0.514 (SA0.3), 0.447 (SA1.0), and 0.397 (SA3.0), demonstrating the model's predictive accuracy.
Overall, these preliminary results show that combining waveform features, network topology, and LLM-informed event metadata can substantially enhance station-level IMs estimation,

achieving better results in challenging conditions (north or south Italy) with respect to previous similar approaches. This method has potential for rapid earthquake characterization and early warning, where timely and accurate ground-motion predictions are essential.

Deep Learning and ray path analysis illuminate spatiotemporal variations in fault zone properties during the seismic cycle

Gabriele Paoletti[1,*], Daniele Trappolini[2], Elisa Tinti[1], Fabio Galasso[2], Cristiano Collettini[1], Chris Marone[1,3]

[1]*Sapienza University of Rome, Department of Earth Sciences, Rome, Italy*
[2]*Sapienza University of Rome, Department of Computer, Control and Management Engineering, Rome, Italy*
[3]*Pennsylvania State University, Department of Geosciences, University Park, PA, USA*

**Corresponding author: gabriele.paoletti@uniroma1.it*

Earthquake fault zones evolve continuously throughout the seismic cycle, yet determining whether such evolution can be directly detected in seismic waveforms remains a fundamental challenge. While previous studies have demonstrated that machine learning models can distinguish earthquakes occurring at different times, it is still unclear whether this capability reflects intrinsic temporal changes in the fault system or is instead modulated by how seismic waves sample structurally complex regions. In particular, the role of ray-path geometry in controlling the observability of time-dependent fault-zone properties has not been systematically investigated.

Here, we address this gap by combining deep learning and ray-path analysis using the 2016–2017 Central Italy seismic sequence as a natural laboratory. We analyze raw three-component seismic waveforms recorded across a dense network and train a Convolutional Neural Network (CNN) to distinguish events occurring in different temporal windows of the seismic cycle. The approach is fully data-driven and does not rely on predefined features or physical parameterization, allowing us to test whether temporal information is inherently encoded in the waveform data.

Our results show that seismic waveforms are consistently distinguishable across time, indicating that temporal evolution of the fault zone leaves a measurable imprint in the wavefield. However, this separability is not spatially uniform. Classification performance in the post-seismic phase varies systematically across stations and is highest when ray paths intersect highly damaged regions of the fault zone, particularly near the fault core. In contrast, stations sampling more external or less structurally perturbed volumes exhibit reduced discriminative capability.

In the pre-seismic phase, temporal separability remains broadly stable across most configurations, but a marked increase in performance is observed at a station sampling the deep nucleation region in the final weeks preceding the mainshock. This localized enhancement suggests that detectable changes in the medium may occur shortly before failure, although such observations remain case-specific and do not imply predictive capability.

These results demonstrate that fault-zone evolution is detectable in seismic waveforms in an operational sense, but that its observability is strongly controlled by propagation geometry. More broadly, this study shows that deep learning applied to raw seismic data can act as a sensitive probe of time-dependent changes in natural fault systems, while also highlighting the need to account for spatial sampling effects.

Unsupervised encoding of seismo-volcanic signals of a full eruptive cycle at Piton de la Fournaise (2014-2023)

Marie A. Gärtner*, Reza Esfahani, Michel Campillo, Nikolai Shapiro

ISTerre, CNRS/UGA, Grenoble, France

**Corresponding author: marie.gartner@univ-grenoble-alpes.fr*

Active volcanoes generate a wide variety of seismic signals that are linked to different physical processes. These seismo-volcanic signals provide valuable insight into the state and dynamics of a volcano's plumbing system. However, their complexity and variability make it challenging to link these signals to underlying physical processes.
Here we present an unsupervised machine learning workflow to extract statistically significant features directly from continuous seismic recordings. Our approach is based on a deep scattering network, a lightweight convolutional neural network with fixed wavelet-based kernels, that provides robust and interpretable representations of complex signals across multiple scales. Since this method does not rely on prior information like catalogs and labels, it enables the exploration of the full complexity of seismograms and allows for a more comprehensive characterization of volcanic activity.
We apply our methodology to a decade of three-component seismic data (2014 to 2024) recorded at Piton de la Fournaise on La Réunion, France, one of the most active volcanoes in the world. Using data from multiple stations of the permanent network, we analyze the most recent complete eruption cycle and identify distinct patterns in the continuous seismic signal that are associated with different phases of volcanic activity. Furthermore, we are able to detect signal patterns that are not captured by the observatories' routine monitoring.

Automated AI agent for earthquake location and magnitude estimation: application to Campi Flegrei caldera 2014-2025 seismicity

Andrea Sollai*, Aldo Zollo, Grazia De Landro, Salvatore Cuomo

University of Naples Federico II, Naples, Italy

**Corresponding author: andrea.sollai@unina.it*

The rapid development of artificial intelligence has led to its increasingly widespread application in seismology, particularly in the automatic detection and phase picking stages of seismic catalogs construction. Despite these advances, the use of machine learning for downstream tasks such as earthquake location and magnitude estimation remains comparatively less explored, especially in fully automated frameworks.

In this study, we propose a novel machine learning–based, region-specific approach for the simultaneous estimation of hypocentral location and magnitude of seismic events. The method is based on a Random Forest model trained on a specifically designed dataset that incorporates intra-event differential arrival-time features. These include single-station S–P time differences (Ts–Tp) and inter-station differential arrival times (ΔTp and ΔTs) computed for station pairs for the same seismic event. For each seismic event, the model is trained to predict both the hypocentral coordinates and the event magnitude, enabling a unified framework for multi-parameter estimation.

A key aspect of the proposed approach is the strategy adopted for magnitude estimation. Since differential arrival times do not directly encode information about event magnitude, the model exploits the availability and spatial distribution of the input features, such as the number of valid observations per event and their spatial distribution across the seismic network.

This results in a data-availability-driven magnitude prediction scheme, which implicitly captures the unknown relationship between event size and network detection threshold.

The method is tested on a high-precision relocated seismic catalogue from the Campi Flegrei caldera (Italy), including approximately 20,000 events recorded between 2014 and 2025, with magnitudes ranging from M_d 0.0 to M_d 4.6. The catalogue was obtained using the high-precision NLL-SSST-WC technique [Lomax et al., 2022] in combination with a high-resolution 3D P- and S-wave velocity model [De Landro et al., 2025]. This dataset provides an ideal benchmark for training and validation, thanks to its high location accuracy and consistency.

The results demonstrate high accuracies in both earthquake location and magnitude estimation. In particular, the predicted hypocentral uncertainties are comparable to those of the high-precision catalogue used for training, while magnitude estimates show good agreement across the full magnitude range. These findings highlight the robustness of the proposed feature representation and its effectiveness in capturing both spatial and observational constraints.

The proposed AI framework naturally incorporates the most advanced relative earthquake location approaches combined with the best updated 3D velocity model but using bulletin phase arrival times at the available stations. It can be integrated into fully automated AI agents for seismic monitoring systems, complementing existing automatic detection and potentially used for early warning pipelines by providing real-time, evolutionary estimates of earthquake location and magnitude.

Quality evaluation of assimilation-based forecast of rate-and-state governed fault analog

Bharath Shanmugasundaram*, Harsha Bhat, Romain Jolivet

École Normale Supérieure - Université PSL, France

**Corresponding author: bharath.shanmugasundaram@ens.psl.eu*

During an earthquake, the frictional resistance of a fault suddenly drops to release the elastic energy that has been accumulating over decades to centuries. In addition to the steady increase of stress on faults due to tectonics, external perturbations have been shown to modulate the fault behavior over a wide range of time scales. The spring block slider model following rate-and-state friction framework with velocity-weakening behavior undergoing periodic perturbations has been known to host complex stick-slip events ranging from fast earthquakes to slow earthquakes, making it a good analog of a simple fault. Accurate characterization of system state and key parameters is critical for understanding the triggering mechanisms and ultimately improving seismic hazard assessment. In this work, we employ ensemble-based data assimilation techniques to carry out state and joint state-parameter estimation in a tidal modulated spring slider. We perform twin experiments to estimate the rheological parameters governing fault friction, leveraging prior knowledge of the tidal perturbation period and amplitude. In this scenario, we compare the iterative ensemble Kalman smoother (I-EnKS) with ensemble Kalman filter (EnKF) variants for joint state-parameter estimation. Using the smoothed estimates, we assess forecast quality by evaluating prediction accuracy over multiple recurrence intervals. To account for model uncertainties, we incorporate additive stochastic forcing to examine its effect on state-parameter estimation and forecast accuracy.

Foundation models reveal seismic damage evolution in fault zones

Sadegh Karimpouli[1,*], Patricia Martínez-Garzón[1,2], Frederik Tilmann[1,3], Georg Dresen[1], Marco Bohnhoff[1,3]

[1]*GFZ Helmholtz Center for Geosciences, Potsdam, Germany*
[2]*RWTH Aachen University, Aachen, Germany*
[3]*Freie Universität Berlin, Department of Earth Sciences, Berlin, Germany*

**Corresponding author: sadegh.karimpouli@gfz.de*

The progressive development of seismic damage strongly influences earthquake rupture processes and thus seismic radiation. However, detecting gradual changes in damage using traditional seismological approaches remains difficult. In this study, we explore whether representations learned by seismic foundation models can capture the temporal evolution of damage through variations in seismic coda waves. We analyze features extracted by the Seismic Language Model (SeisLM), a large self-supervised model trained on extensive collections of seismic waveforms, and quantify their temporal variability using a distance-based dissimilarity measure.

To evaluate the approach, we generate synthetic datasets using two-dimensional numerical simulations with progressively evolving damage properties of the fault zone, including increased fracture density and attenuation together with reduced seismic velocities and density. The computed dissimilarity metrics show a systematic increase as damage intensifies and remain relatively stable under variations in source and receiver geometry.

We then apply the same analysis framework to seismic waveform recordings from the 2016 Amatrice–Visso–Norcia earthquake sequence in central Italy. The results reveal consistent temporal patterns in dissimilarity across events sampling with different spatial volumes, from local to more regional scales. Using the mean feature vector of the first 100 events as a baseline, dissimilarity values drop immediately after major earthquakes and subsequently rise over time. Independent observations from P-wave velocity changes derived from seismic tomography support these patterns and suggest enhanced damage accumulation followed the large events in the sequence, and a subsequent partial healing of the media.

Overall, our results indicate that waveform representations learned by seismic foundation models provide a sensitive and integrative indicator of evolving crustal damage and near-surface features, offering a data-driven complement to conventional seismological monitoring techniques.

Machine Learning based prediction of laboratory fault peak slip-rate trends using seismicity catalog

Xu Han[1,*], Sadegh Karimpouli[1], Marco Bohnhoff[1,2], Georg Dresen[1]

[1]*GFZ Helmholtz Centre for Geosciences, Potsdam, Germany*
[2]*Freie Universität Berlin, Department of Earth Sciences, Berlin, Germany*

**Corresponding author: haha1820798358@gmail.com*

Fluid injection into subsurface reservoirs can facilitate the reactivation of natural faults, inducing slip at varying rates depending on the in-situ stress state, fluid pressure, and hydraulic properties. As a laboratory analogue of injection-induced earthquakes, we conduct four sets of fluid-driven fault slip experiments on cylindrical Bentheim sandstone samples containing pre-cut faults, subjected to different fluid injection rates and initial confining pressures. Across these experiments, we observe two distinct behaviors in the evolution of peak slip rate: it either increases or remains approximately constant from cycle to cycle. We attribute these contrasting trends to the combined effects of elevated pre-stress, pore pressure evolution, and the development of a localized gouge layer. As these trends are related to fault stability, they may serve as important indicators for adapting subsequent fluid injection strategies. To assess whether these trends are predictable using seismicity patterns, we investigate whether the next peak slip rate will increase or remain constant relative to the previous peak slip rate. To this end, we train a Convolutional Neural Network (CNN) classifier with acoustic emission (AE) catalogs. The classifier utilizes features derived from sequences of 100 consecutive AE events (e.g., magnitude, inter-event time, and three-dimensional location) within any time window of a given injection cycle to classify the peak slip-rate trend. Results reveal that the model achieves a balanced classification accuracy of 96%, indicating that the peak slip-rate trend can be predicted with high reliability. This achievement suggests that the evolution of fault slip rates may be preceded by detectable and potentially forecastable seismicity patterns, at least for the two trends observed in our experiments. This framework may potentially be extended to enhanced geothermal systems and contribute to the improvement of the warning protocols in enhanced geothermal systems.

Masked graph neural network for rapid ground motion prediction in Italy

Ilaria Oliveti[1,*], Daniele Trappolini[2], Licia Faenza[3], Dario Jozinović[4], Alberto Michelini[1]

[1]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Osservatorio Nazionale Terremoti, Rome, Italy*
[2]*Sapienza University of Rome, Rome, Italy*
[3]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Bologna, Bologna, Italy*
[4]*Swiss Seismological Service (SED), ETH Zurich, Zurich, Switzerland*

**Corresponding author: ilaria.oliveti@ingv.it*

We propose the application of TISER-GCN [Bloemheuvel et al., 2022] to the Italian seismic network. TISER-GCN is a graph neural network (GCN) architecture designed to predict the maximum intensity measurements of ground shaking (PGA, PGV, SA(T)) from the pattern of the recorded seismic waveforms in the very first seconds after the earthquake's first P-wave arrival times. The input data used in this study consist of the three-component earthquake waveforms recorded primarily by the Italian National Seismic Network and collected in INSTANCE - a seismological benchmark dataset. We have selected about 3,000 earthquakes between 2005 and 2020 with magnitude ≥ 3.0 and depth ≤ 50 km. To gather a high-quality dataset for training TISER-GCN, we have adopted various criteria to avoid possible signal saturation and to privilege the use of true acceleration recordings whenever possible. TISER-GCN adopts an innovative approach to incorporating both the temporal and spatial dimensions of data through graph structures, leading to significant advancements and positioning it as one of the state-of-the-art models for this type of task. Noteworthy, it does not require previous knowledge of the earthquake source (location and magnitude) and can provide accurate predictions of ground motion at farther seismic stations from the initial recordings of the earthquake at the nearest stations. The challenge is applying the model to the entire national network, comprising nearly 600 stations, thereby requiring to overcome technical issues. This expansion requires adaptation to a sparse graph with null values for many stations, a significant step forward from previous works restricted to a smaller area featuring only 39 stations. Results show that masking stations where the P-wave arrives within the first 10 seconds, combined with the integration of additional information, reduces the mean squared error (MSE) by up to 6% for peak ground acceleration (PGA) and 5.5% for peak ground velocity (PGV), compared to the unmasked baseline. Moreover, the proposed approach yields near-zero median residuals across all IMs, mitigating the systematic underestimation observed when using a ground motion model specifically developed for Italy.

Furthermore, incorporating pre-trained deep learning phase-picker weights, such as those from the PhaseNet encoder, yields noticeably better performance, confirming the benefits of improved generalization. These findings indicate that the model provides accurate predictions of ground motions, comparable to those obtained with the original TISER-GCN, which, however, requires a fixed seismic network geometry.

AI-enabled digital twins for volcano monitoring: insights from the DT-GEO project

Flavio Cannavò[1,*], Paolo Papale[2], Deepak Garg[2], Martina Allegra[1], Rebecca Bruni[2], Erica De Paolo[1], Chiara Montagna[2], Gilda Currenti[1], Danilo Reitano[1], Michele Carpené[3], Piero Lanucara[4]

[1]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Catania, Italy*
[2]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Pisa, Pisa, Italy*
[3]*CINECA, Casalecchio di Reno (BO), Italy*
[4]*CINECA, Rome, Italy*

**Corresponding author: flavio.cannavo@ingv.it*

The development of DTs for geophysical systems is a key step in gaining insight into dynamic natural processes through data in real-time. In the framework of the DT-GEO project, AI, HPC, and different kinds of observational data were used to develop a scalable DT for the study of volcanic unrest phenomena. Our approach is based on the use of real-time data obtained from observation networks together with data-driven AI models, which were trained using a large-scale numerical simulation of physical processes responsible for the data signatures.

The DT model is designed with a modular workflow that comprises multiple components, including the real-time acquisition of data and its pre-processing, a machine learning based algorithm for the detection and interpretation of unrest signals, and inversion algorithms based on an artificial intelligence a physics-based numerical model that relates surface measurements to subsurface phenomena. The key element in this strategy is the generation of massive synthetic datasets through high-fidelity numerical simulations. This allows for the training of artificial intelligence models to obtain rapid solutions to inversion problems that would otherwise require substantial computing resources. Using this strategy, the model can deliver almost real-time probabilistic scenarios related to volcanic unrest, thus enhancing situational awareness.

Although demonstrated on a specific volcanic system, the methodology is intentionally generic and adaptable to a wide range of geophysical contexts. The reliance on scalable HPC infrastructures and open, reproducible workflows ensures that the approach can be extended to other volcanoes and potentially to different natural hazards characterized by complex, multi-scale dynamics.

Future work will aim to make models better at handling different geological conditions and bring in more data sources like seismic and remote sensing. There will also be a push to include more advanced ways to measure uncertainty. Moving toward fully linked, multi-physics Digital Twins should help improve how well predictions work. At the same time, there will be a focus on making these systems practical to use, improving their reliability, speeding up computations, and fitting them into tools that help with hazard assessment and managing risks.

The DT-GEO experience demonstrates the potential of AI-enabled Digital Twins to bridge observational data, physical modelling and AI, paving the way for next-generation monitoring and forecasting tools in volcanology and beyond.

Application of data science and Machine Learning techniques for the prediction of induced seismicity

Leticia Garay Romero[1,2,*], Licia Faenza[2], Alexander Garcia[1], Anna Maria Lombardi[2]

[1]*University of Bologna (UNIBO), Bologna, Italy*
[2]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Bologna, Bologna, Italy*

**Corresponding author: leticia.garayromero2@unibo.it*

The prediction of induced seismicity is a critical challenge for geological risk management and the safe operation of industrial facilities, such as geothermal projects. This study focuses on the Cooper Basin in Australia. We applied data science and artificial intelligence techniques to analyze seismic time series, integrating two data sources: discrete seismological events (23,285 events) and continuous operational data sampled every 2 minutes (33,839 records).
The main objective was to develop machine learning models to predict, in future time windows of 10, 30, 60, and 90 minutes, two key variables: the number of seismic events or the maximum magnitude. The XGBoost and Random Forest algorithms were trained and compared. Model performance was evaluated using the R^2, RMSE, and MAE metrics, and their interpretability was analyzed using SHapley Additive exPlanations (SHAP).
The results demonstrate that both models generate predictions consistent with the observations, showing better predictive performance in the longer time windows (60 and 90 minutes).
This approach provides a valuable framework for the monitoring and proactive risk assessment of geothermal operations.

Deep Learning-based urban acoustic source identification from DAS data in the time-space-frequency domain

Yichen Zhong[1,*], Chen Gu[2], Michael Fehler[2], German A. Prieto[3], Peng Wu[1], Zhi Yuan[1], Jiawei He[1], Chenrui Fan[1], Yufeng He[1]

[1]*Tsinghua University, Beijing, China*
[2]*Massachusetts Institute of Technology, Cambridge, MA, USA*
[3]*Universidad Nacional de Colombia, Bogotá, Colombia*

**Corresponding author: zhongyc22@mails.tsinghua.edu.cn*

A key step toward improving disaster prevention and mitigation in sustainable smart cities is to enhance the sensing capability of existing communication infrastructure, thereby providing reliable and continuous information for urban management during emergencies. Distributed Acoustic Sensing (DAS) is a promising technology for this purpose because it transforms existing fiber-optic cables into dense sensing arrays with wide spatial coverage, high temporal resolution, and broad applicability in urban environmental perception, structural health monitoring, and hazard warning. In this study, field data were collected using a DAS array deployed along the edge of the Guye area in Tangshan, China, where the 1976 Tangshan earthquake occurred. The recorded signals contain vibrations generated by both natural and anthropogenic acoustic sources, providing a valuable dataset for urban acoustic source analysis in a seismically significant region.

We first investigated the characteristics of different source types in the temporal, spatial, and frequency domains. To better capture the multidimensional features of DAS signals, the raw data were transformed into a time-space-frequency representation, from which space-frequency spectrograms were constructed for acoustic event detection and classification. This representation propagation characteristics, and spectral signatures of different sources, making it particularly suitable for distinguishing complex urban vibrations. Based on these observations, we developed a deep learning-based framework for the detection and classification of multiple acoustic sources, including earthquakes, vibrator vehicles, traffic flow, and industrial activities. To support model training, a dedicated dataset was constructed from the acquired DAS field measurements and annotated using the label-transfer method proposed in this study, which improves labeling efficiency and facilitates the extraction of representative acoustic events in the space-frequency domain.

Experimental results demonstrate that the proposed method achieves an overall classification precision of 98.6% and an mAP50–95 of 90.0%, indicating strong performance in both recognition accuracy and localization quality. In addition, the framework processes each DAS spectrogram frame in 6.9 ms on an NVIDIA GeForce RTX 4090 GPU, demonstrating its capability for real-time urban monitoring. These results show that the proposed source-identification scheme can reliably monitor routine urban activities over long distances while maintaining high accuracy and sensitivity to abnormal events.

Beyond the current application, the framework is scalable and can be extended to recognize additional source classes or adapted to other specialized DAS datasets. Overall, this study highlights the significant value of DAS for enhancing urban environmental perception and hazard information analysis in smart cities. The proposed method also shows strong potential for broader applications, including earthquake detection, site-effect studies, and anomaly detection in complex urban environments.

Explainable physics-guided feature extraction from time-lapse ultrasonic waveforms for predicting laboratory earthquakes

Prabhav Borate*, Jacques Rivière, Daniel Kifer, Parisa Shokouhi

The Pennsylvania State University, University Park, 16802, USA

**Corresponding author: pxs990@psu.edu*

Physics-based handcrafted features such as P-wave amplitude derived from the early parts of time-lapse ultrasonic monitoring waveforms, are shown to successfully predict impending laboratory earthquakes. This study investigates whether feature extraction can be automated by leveraging the full ultrasonic waveforms. We compare four different models: a handcrafted feature-driven model (baseline), a convolutional neural network (CNN) model, an explainable data-driven (XDD) model, and a model with an explainable physics-guided (XPG) feature. The baseline model uses the spectral amplitude of the P-wave packet calculated at the center frequency as the feature while CNN relies on CNN and pooling layers for feature extraction. The XDD model combines CNN feature extraction with long short-term memory units and an attention layer to capture the time history of the features as well as their significance. The novel XPG extracts a generalized spectral amplitude feature through a custom multilayer perceptron model designed to learn the linear combination of signals' spectral amplitudes (physics-guided feature) without prescribing specific frequencies. Although CNN-extracted features resemble handcrafted ones, both the CNN and XDD models underperform compared to the baseline model while the XPG model outperforms all the other models. While the attention mechanism in the XDD model highlights the importance of the first P-wave packet, the XPG model also identifies the later-arriving shear wave packet as significant. The improvement in prediction with the inclusion of shear waves is physically intuitive; S-waves have particle motion polarized in the direction of fault shear and are likely more sensitive to the frictional state than P-waves. This work demonstrates the value of using domain knowledge in automatic feature extraction with potential applications in seismic data analysis.

Beyond prospectivity maps: multimodal geospatial AI for critical mineral systems and geothermal prospectivity exploration

Velimir Vesselinov

EnviTrace LLC, USA

Corresponding author: monty@envitrace.com

The global energy transition is reshaping the demand for subsurface intelligence. Critical minerals and geothermal resources are both essential for future economic development. Yet their discovery and extraction remain limited by sparse observations, fragmented datasets, uncertain geologic processes, and high exploration costs. Traditional prospectivity mapping has helped organize evidence, but too often it stops at static maps rather than producing adaptive, uncertainty-aware decisions.

This talk presents our work towards the development of a multimodal geospatial AI framework for moving beyond conventional prospectivity maps toward integrated, decision-ready exploration systems. The approach combines geologic maps, geochemistry, geophysics, remote sensing, structural data, thermal indicators, borehole information, and expert knowledge within machine learning workflows designed to account for data gaps, scarce labels, and heterogeneous spatial data.

The presentation will compare our novel approaches with traditional methods such as gradient boosting and geostatistical techniques.

We conclude that geospatial AI should not replace geologists, geophysicists, or geothermal experts. It should be applied as an important tool to sharpen the analyses, expose hidden patterns, quantify uncertainty, and help prioritize exploration decisions under real-world constraints. The future of resource discovery is not just better maps, but better reasoning systems built from data, physics, and domain expertise.

Time-reversal neural operator for earthquake rupture inversion via data fusion of broadband seismic arrays and submarine DAS

Chen Gu[1,2,*], Zhuoyu Chen[1], Jiawei He[1], Yufeng He[1], Michael Fehler[2], Yichen Zhong[1], Peng Wu[1], Zhi Yuan[1], Chenrui Fan[1], Shigui Zhang[1], Jiachao Xu[1], Xiaoyong Qin[1], Xiaoyan Fan[1], Hongbo Sun[1]

[1]*Tsinghua University, Beijing, China*
[2]*Massachusetts Institute of Technology, Cambridge, MA, USA*

**Corresponding author: guchch@mit.edu*

Distributed Acoustic Sensing (DAS) is rapidly reshaping observational seismology by enabling dense, large-aperture measurements along fiber-optic cables, particularly in offshore regions where conventional seismic coverage remains sparse. Yet, a fundamental gap persists: broadband seismic arrays provide stable, large-scale constraints, while DAS arrays offer high-resolution strain-rate measurements over limited spatial extents. Existing inversion frameworks rarely exploit the complementary nature of these two modalities in a unified, physics-consistent manner.

We propose a Time-Reversal Neural Operator (TRNO) for earthquake source rupture inversion that integrates large-scale broadband seismic arrays with small-aperture, high-resolution submarine DAS arrays. The method builds on operator learning to approximate the solution operator of wave propagation, while embedding a recurrent time-reversal mechanism that iteratively reconstructs the spatio-temporal evolution of rupture processes. By combining neural operator architectures with time-reversal dynamics, the framework naturally accommodates multi-scale data fusion and enforces physical consistency across heterogeneous observations.

We apply the proposed approach to a unique dataset consisting of a 525 km submarine fiber-optic cable system connecting Hainan and Hong Kong, together with regional broadband seismic stations from the China Array. The DAS system has continuously recorded strain-rate data for over two months since December 24, 2025, capturing both ambient noise and regional seismicity. Synthetic experiments based on realistic fault geometries along nearby convergent margins, including the Manila Trench, the Ryukyu–Taiwan plate-boundary system, and the Indo-Burma subduction zone, demonstrate that TRNO robustly recovers rupture kinematics under complex velocity structures and limited offshore coverage. We further validate the method using the M_W 6.6 Yilan Earthquake 2025, where the joint inversion of broadband and DAS data reveals detailed rupture evolution consistent with independent geophysical constraints, while significantly improving spatial resolution near the offshore fault segment.

This work highlights the potential of neural operator-based time-reversal methods for next-generation earthquake rupture inversion and demonstrates a practical pathway toward integrating heterogeneous sensing modalities for offshore earthquake physics.

Mitigating atmospheric noise in InSAR time series over volcanoes using a convolutional neural network

Rebecca Bussard*, Christelle Wauthier

The Pennsylvania State University, University Park, 16802, USA

**Corresponding author: rqb5400@psu.edu*

Ground deformation can be a key precursor to volcanic activity, and can also be used to constrain magmatic source parameters such as depth and volume change. Interferometric Synthetic Aperture Radar (InSAR) measures centimeter-scale ground deformation through time across volcanic regions, but atmosphere over these regions can introduce noise into the data, making it hard to constrain ground movement related to volcanic sources. Since atmospheric conditions are affected by regional climate, local topography, and time of day, implementing a general correction method on InSAR time series across numerous volcanoes is challenging. To address this challenge, this work utilizes a deep learning, convolutional neural network (CNN) trained on thousands of synthetic unwrapped InSAR time series to isolate ground deformation from atmospheric noise. Each synthetic time series contains deformation from 0-2 sources with source types including point, spheroid, sill, and fault. Atmospheric noise is synthetically introduced two ways in separate trainings: 1) utilizing topographic data over twenty volcanoes of varying types to calculate stratified phase and combining this stratified phase with spatially correlated noise, and 2) using phase maps from the Generic Atmospheric Correction Service (GACOS) over the same group of twenty volcanoes. The trained CNNs are tested on synthetic time series with signal to noise ratios (SNRs) ranging 0.1-5.0 to retrieve below centimeter scale error for most SNR values. Real InSAR data spanning 2015-2016 over Masaya volcano in Nicaragua is input into the models, with predicted deformation of Masaya's point source isolated from atmospheric noise and in agreement with available GNSS deformation data.

In-context learning enables global geothermal temperature prediction from sparse local observations

Daniel O'Malley*, David Coblentz, Nathan Debardeleben, Chris Johnson, John Kath, Pablo Lara, Earl Lawrence, Sandro Malusa, Arnab Mazumder, Mohamed Mehana, Javier Santos, Bharat Srikishan, Hari Viswanathan

Los Alamos National Laboratory, USA

**Corresponding author: omalled@lanl.gov*

Subsurface data is often sparse, raising a central question for AI in the geosciences: when trained only on borehole temperatures, can a model recover a broader, physically meaningful representation of the subsurface, or does it simply interpolate nearby measurements? Here we introduce Transparent Temperature, a transformer-based, context-aware framework that infers continuous temperature-at-depth fields from sparse borehole observations while producing calibrated uncertainty. Interpretability analyses show that, despite never being trained on them, the model learns internal representations of subsurface properties including seismic velocities, geochemistry, and crustal structure, and uses these representations in physically consistent ways. In parallel, the model delivers strong predictive performance: in the contiguous United States it achieves an MAE of 4.8 °C (25% lower than the physics-informed Stanford Thermal Model and 20% lower than a model based on AlphaEarth embeddings) with well-calibrated probabilistic predictions (Kolmogorov-Smirnov statistic is 2.5%). With frozen weights, it transfers to Australia, Canada, and the United Kingdom with accuracy comparable to the United States using only 20 observations from those regions to make a prediction.
The approach also preserves sharper thermal gradients caused by, e.g., hydrothermal flows than prior techniques. We perform a series of ablation tests to identify the ingredients needed for accurate subsurface modeling with transformer-based techniques. These ingredients enable the model to learn a multifaceted representation of the subsurface that extends beyond temperature to include geophysical and geochemical properties and to recover known physical relationships (e.g., increasing the model's representation radioactive element concentrations increases predicted temperature) despite no explicit training objective to incentivize this behavior.

The Transparent Earth: a multimodal foundation model for Earth sciences

Sandro Malusà*, Javier Estrada Santos, Pablo Eduardo Espinoza Lara, Christopher Johnson, David Coblentz, Daniel O'Malley

Los Alamos National Laboratory, USA

**Corresponding author: smalusa@lanl.gov*

Predicting subsurface properties where direct measurements are unavailable or prohibitively expensive remains a fundamental challenge in Earth sciences. At the same time, large amounts of heterogeneous measurements of different geophysical quantities are widely available in public datasets. Each offers only a partial and often limited view of the subsurface due to data scarcity, noisy measurements, sparse measurements and/or sparse coverage with data available only in restricted specific regions of the globe. Recent advances in machine learning provide us with new modeling strategies to tackle this problem.

In this work, we present The Transparent Earth, a monolithic multimodal foundation model designed to process data from measurements of different geophysical quantities (including thermal, stress, earthquakes and magnetic information, to name a few) and use them jointly to provide a global representation of the Earth as a unified system. The model is built on a transformer-based backbone that is agnostic to both spatial position and input resolution. The model is trained with a large set of measurements from different quantities (modalities) and learns to reconstruct each of them given a set of multi-modal observations in the surrounding of the query location. Therefore, the model learns complex cross-modal relationships across heterogeneous geophysical data that generalize to regions unseen during training. This formulation enables a key capability: at inference time, the model can reconstruct a target field regardless of the specific type of observations available in its vicinity. Therefore, in contrast to unimodal approaches, The Transparent Earth can infer a given modality even where that same quantity has never been measured locally, provided that other informative geophysical observations are available. This makes the model valuable for scientific discovery and exploration of regions where specific (typically expensive) measurements are not available. Together, these features position The Transparent Earth as a "next generation" foundation model for making subsurface predictions in the geosciences, with the potential to support multimodal interpolation, cross-modal inference, and data-driven exploration at global scale.

Scale-aware seismic phase arrival detection via continuous dilatable convolutional kernels

Sacha Lapins

University of Oxford, UK

Corresponding author: sacha.lapins@earth.ox.ac.uk

Deep learning seismic phase arrival detection models increasingly underpin earthquake monitoring and catalogues, yet current architectures share a fundamental limitation. At inference time, they comprise a fixed set of convolutional filter parameters, optimised for the specific signal characteristics and sampling rate of their training data. Applying these models to data with differing characteristics (e.g., from different source types, hypocentral distances, magnitude ranges, or sensor modalities) typically results in degraded performance. Simply expanding training sets to encompass more signal, noise or event diversity is unlikely to fully resolve this issue of generalisation. Seismic signal characteristics (e.g., amplitude, frequency content, duration) are known to scale with source process size/type, source-receiver distance and instrument sampling rate, and fixed-parameter models have finite capacity to accommodate this diversity (which is often exhibited over exponential orders of scale). Previous approaches have addressed this challenge through data-side strategies: e.g., resampling/filtering input data, and aggregating predictions across multiple preprocessing scales and models, to make new data resemble the training data of pre-trained models. Here I take the opposite approach: rather than adapting data to the model, I explore whether the model can be made to inherently adapt to the data.

I present a scale-aware convolutional architecture for seismic phase arrival detection built on the idea of continuous kernel parameterisation, which I term CODA (COntinuous Dilated-scale Aggregation). Rather than learning discrete filter weights at a fixed resolution/scale, convolutional kernels are represented by continuous functions that can be rediscretised on-the-fly at arbitrary temporal scales during both training and inference, allowing kernels to effectively dilate and contract across the model's full feature-extraction parameter space. This design is motivated by the empirical observation that first-layer kernels learned by seismological convolutional networks (e.g., GPD, PhaseNet) resemble wavelet-like basis functions, which seismologists have long used for time-scale decomposition of seismic waveforms. Making kernel dilation explicit, learnable, and ubiquitous across feature extraction stages therefore provides a natural architectural formulation for seismic waveform data.

Within a CODA block, a scale integration module weights contributions across a range of effective dilations (scales) for each model feature, producing representations that adapt to the signal characteristics of each input waveform. The explicit use of kernel scaling weights carries interpretative value beyond generalisation, providing an inherent coordinate system for post-detection clustering, classification, or filtering of arrivals by scalable signal attributes, such as frequency content, P-S timing, or event duration.

In this presentation, I introduce the CODA architecture and present experimental results from training and evaluation across diverse benchmark datasets available through SeisBench, spanning a range of monitoring settings, source types, sampling rates, and magnitudes. Through these results, I assess whether continuous, scale-aware kernel parameterisation improves robustness across heterogeneous seismic datasets and reduces reliance on resampling, filtering, or retraining based adaptation, and discuss some future directions for this approach.

Towards DL Phase Picking for diverse volcanic seismicity: a case study at Tungurahua

Abigail Robinson[1,*], Prof. James Hammond[2], Dr Sacha Lapins[3], Dr Alex Steele[1]

[1]*University College London (UCL), Department of Earth Sciences, London, UK*
[2]*University of London, Birkbeck, London, UK*
[3]*University of Oxford, Department of Earth Sciences, Oxford, UK*

**Corresponding author: abigail.robinson.23@ucl.ac.uk*

Deep learning phase-picking (DL-PP) models are transforming volcano seismology, enabling the construction of more complete seismic catalogues and higher resolution subsurface imaging of systems including Kilauea [Wilding et al., 2023], Yellowstone [Armstrong et al., 2023], Santorini [Lomax et al., 2025], and Campi Flegrei [Tan et al., 2025].

Tungurahua, an andesitic stratovolcano in Ecuador that experienced major unrest between 1999 and 2016, is a compelling target for this approach. The IG-EPN has operated a broadband seismic network at Tungurahua since 2006 and created an extensive manually labelled catalogue of ~299,000 events spanning over a decade of diverse activity and signal types, providing an exceptional foundation for DL-PP development and evaluation. At the same time, most events have been classified at only a single reference station, and far fewer have been located, leaving many questions about subsurface structure and processes during this eruptive period unanswered. A DL-PP enhanced catalogue creation workflow therefore holds potential to greatly enhance the number and quality of event locations made at Tungurahua, and in turn, to generate a clearer picture of subsurface structure and processes.

However, applying existing DL-PP models at Tungurahua poses significant challenges. Most models have been developed and evaluated in settings dominated by tectonic and volcano-tectonic (VT) seismicity, and recent work has shown that the performance of most DL-PPs degrades systematically with decreasing frequency content [Zhong and Tan, 2024]. Tungurahua's seismicity is exceptionally diverse, spanning VT earthquakes, long-period (LP) events, very-long-period signals, tornillos, explosions, and multiple tremor types, and the extent to which existing models can handle this diversity is unknown. Before deploying DL-PP models to enhance the catalogue, it is therefore essential to understand where they succeed, where they fail, and what can be done to improve them.

This poster presents a systematic evaluation of existing DL-PP models applied to the Tungurahua dataset. I assess model performance across Tungurahua's diverse signal types and examine newly detected events not present in the manual catalogue. These results provide a foundation for selecting and improving DL-PP tools for Tungurahua, and offer broader insights into the applicability of existing models at volcanoes where non-VT seismicity dominates.

High-resolution lithospheric thickness model of the Southern Indian Shield using Physics-Informed Neural Networks and geophysical data integration

Arka Roy*, Jörg Ebbing

Institute of Geosciences, Kiel University, 24118 Kiel, Germany

**Corresponding author: arka.phy@gmail.com*

We present a high-resolution model of lithospheric thickness for the Southern Indian Shield, developed from geoid, gravity gradient, and elevation data using Physics-Informed Neural Network (PINN) framework to enable faster and more accurate estimation of lithospheric structure. The approach is based on the principle of local isostasy, which assumes that the lithosphere is in hydrostatic equilibrium and floats on the underlying asthenosphere. A comprehensive training dataset, comprising crustal and lithospheric thicknesses, densities, elevation, geoid height, and gravity gradients is used to train the PINN. The network employs a hybrid loss function that includes (1) data loss, which measures the misfit between model predictions and observed values, and (2) a physics-based component that imposes the isostatic equilibrium constraint on the modelled lithospheric parameters. Integrating these physical constraints allows the PINN to estimate the lithospheric structure both rapidly and accurately. A feed-forward neural network architecture with four input parameters (mean crustal density, gravity gradient, geoid height, and elevation) and two output parameters (lithosphere-asthenosphere boundary [LAB] depth and Mohorovičić discontinuity [Moho] depth) is trained with incorporated isostatic equilibrium constraints. Mean crustal density, treated as the primary unknown parameter, is systematically varied to generate different LAB thickness predictions. The optimal range of mean crustal density is constrained using discrete LAB depths derived from seismic receiver function analysis. Temperature at the Moho, an essential parameter for training data preparation and depth-dependent density estimation of the lithospheric mantle, is derived from the thermal conductivity equation utilizing surface heat flow, radiogenic heat production, and lithospheric thermal conductivity. In this study, surface heat flow is constrained from borehole data, while LAB depth uncertainties are quantified through propagation of errors in thermal conductivity and radiogenic heat production parameters. A comprehensive uncertainty quantification has been performed to develop a high-resolution detailed model of the LAB architecture for the Southern Indian Shield. The Southern Indian Shield consists of several distinct crustal blocks and Archean cratons. Our model effectively delineates craton boundaries, identifies lithospheric roots beneath individual cratons, and provides new insights into the tectonic settings and geodynamic evolution through an efficient deep learning framework.

Advanced data model for the multidisciplinary cataloguing of volcanic ash

Simona Caruso*, Flavio Cannavò, Martina Allegra, Lucia Cacciola, Daniele Andronico, Antonio Cristaldi

INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Catania, Italy

**Corresponding author: simona.caruso@ingv.it*

The study of volcanic ash morphology and componentry is fundamental for understanding the dynamics of explosive eruptions. This work presents a framework for the quickly cataloguing of data related to Etna and Stromboli volcanoes within the INGV-OE surveillance activity.
Our approach combines quantitative morphometry with qualitative textural analysis carried out in the INGV-OE sedimentology laboratory. Using a CAMSIZER® optical analyser, we calculated granulometric and shape data in ash samples to extract specific geometric parameters, including mean size, median, sorting, sphericity, aspect ratio, symmetry, and convexity. Cross-referencing these metrics makes it possible to identify specific "morphological domains" directly related to the eruption intensity and style. Concurrently, observations conducted under a stereomicroscope (Leica M205A) allow for the textural characterization of the sample by identifying three main classes of particles: juvenile material, lithics and free crystals. Juvenile particles are further characterized based on lustre, colour, internal texture, external morphology, and shape (e.g. fluidal, spongy, or blocky). The integration of these two dataset enhances our ability to interpret syn-eruptive conduit processes and strengthens the monitoring of explosive activity.
However, cataloguing large numbers of samples is often a laborious and time-consuming process. This can delay the generation of timely scientific insights, essential for the integration with data from other monitoring networks to effectively interpret the evolution of ongoing eruptive phenomena. This is where deep learning comes in, allowing us to overcome this bottleneck and analyze a much larger volume of data in less time.
With the aim of assisting and speeding up the operator's work, we exploit the acquired high-definition images of the clast, combined with rigorous quantitative analysis performed, to train an AI model for the automatic detection and classification of erupted particles. The proposed Deep Learning framework involves an initial segmentation block, dedicated to the separation of multiple samples in a single image. Subsequently, the extracted particles are classified into the main ash categories.
The computational efficiency of the proposed approach makes it particularly well-suited to the multidisciplinary cataloguing of huge amounts of erupted samples and, as a consequence, for volcano monitoring, thereby enhancing risk assessment.

AI-driven automated and scalable processing of DAS data for seismic monitoring

Miriana Corsaro[1,2,*], Flavio Cannavò[2], Gilda Currenti[2], Simone Palazzo[1], Martina Allegra[2], Philippe Jousset[3], Michele Prestifilippo[2], Concetto Spampinato[1]

[1]*DIEEI, UNICT, Italy*
[2]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Catania, Italy*
[3]*GFZ, Germany*

**Corresponding author: miriana.corsaro@ingv.it*

Artificial Intelligence (AI) is playing an increasingly central role in seismic monitoring, enabling the automatic handling and interpretation of large, complex datasets. Through advanced deep learning methods, AI supports faster and more efficient analysis workflows, while still ensuring coherence with the physical mechanisms governing seismic phenomena.
In this study, we present a fully automated, AI-driven processing pipeline in which distinct deep learning tasks are organized into a sequential and modular chain, guiding the transformation of raw Distributed Acoustic Sensing (DAS) data into reliable earthquake locations. Each stage of the workflow is addressed by a dedicated model, specifically designed to solve a well-defined problem while contributing to the overall system performance.
The first step focuses on seismic phase picking, where deep learning algorithms are used to identify coherent wave arrivals across DAS channels, exploiting both temporal and spatial patterns in the data. The resulting picks are then passed to a second module, where probabilistic machine learning methods perform event association, grouping consistent arrivals across channels into candidate seismic events. Finally, a transformer-based model processes the associated picks to estimate source parameters, including hypocentral location, origin time, and magnitude.
This modular AI pipeline preserves the logical structure of traditional seismological workflows while significantly enhancing each component through data-driven approaches. Applied to the Campi Flegrei area, it demonstrates how combining specialized AI models in a chained architecture enables scalable, robust, and near–real-time seismic monitoring. This approach highlights the potential of AI not only to improve individual tasks, but to integrate them into efficient end-to-end systems for the analysis of complex volcanic environments.

Deciphering pre-eruptive dynamics in Antarctic and Italian volcanoes through Machine Learning

Mónica Ágreda-López[1,*], Irene Rocchi[2], Pier Paolo Giacomoni[2], Maurizio Petrelli[1], Matteo Masotta[2], Sergio Rocchi[2]

[1]*University of Perugia, Department of Physics and Geology, Perugia, Italy*
[2]*University of Pisa, Department of Earth Sciences, Pisa, Italy*

**Corresponding author: monica.agredalopez@gmail.com*

Recent advances in analytical techniques and experimental petrology have led to the generation of increasingly high-dimensional datasets, including crystal-scale chemical maps and multimodal geochemical and textural data. While these datasets offer unprecedented resolution on magmatic processes, they also introduce challenges related to data integration, scalability, and interpretability.

Machine learning (ML) offers a means to extract structure from such datasets. In this study, we explore the opportunities, limitations, and epistemological implications of ML applications in igneous petrology, specifically focusing on pre-eruptive dynamics. Issues of interpretability, calibration, and reproducibility are particularly critical in the context of volcanic hazard assessment, where model outputs may influence decision-making. Here, we discuss these challenges and their implications for the use of ML in volcanology.

We then present a framework that integrates ML with petrological constraints to reconstruct magma storage conditions and plumbing system architectures. The approach combines supervised thermobarometric models with unsupervised clustering applied to major and trace element compositions and high-resolution chemical maps of mineral phases, allowing the identification of distinct crystal populations, quantification of crystallisation conditions, and reconstruction of magma storage systems.

Applications to Antarctic and Italian volcanic systems show that ML can help to identify relevant structures and processes from complex datasets and provide insights into magma reservoir dynamics and pre-eruptive evolution.

Towards Machine Learning-based characterization of seismo-volcanic events in a closed-conduit volcano: the case of Mt. Melbourne, Antarctica

Graziano Larocca[1,*], Danilo Contrafatto[1], Gaetano Giudice[1], Francesco Ciancitto[1], Daniele Pellegrino[1], Alessandro Bonforte[1], Salvatore Consoli[1], Francesco Panzera[2], Mariangela Sciotto[1], Andrea Cannata[1,2], Flavio Cannavò[1], Loredana Contrafatto[3], Simona Scollo[1]

[1]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Italy*
[2]*University of Catania - Biological, Geological and Environmental Sciences Department, Catania, Italy*
[3]*University of Catania - Civil Engineering and Architecture Department, Catania, Italy*

**Corresponding author: graziano.larocca@ingv.it*

Mt. Melbourne, a prominent closed-conduit stratovolcano in Victoria Land (Antarctica), represents a key site for understanding volcanic processes in extreme environments. Although currently characterized by fumarolic activity, its geological history confirms it is an active system, with the latest known eruption dated between the late 18th and early 19th centuries. This research analyzes seismic signals across multiple Antarctic campaigns, spanning several years to the present , integrated within two major frameworks: ICE-VOLC (MultiparametrIC Experiment at antarctica VOLCanoes: data from volcano and cryosphere-ocean-atmosphere dynamics), a multidisciplinary project focused on the internal dynamics of Antarctic volcanoes, and the I-VOLCAN observatory (Italian VOLCanological Antarctic Network), aimed at developing a robust monitoring system, with a primary focus on Mt. Melbourne and a secondary application to Mt. Rittmann.

In this study, we use data from broadband seismic stations located at the volcano's summit. Monitoring Mt. Melbourne presents significant logistical and technical challenges; the harsh Antarctic environment—characterized by prohibitive temperatures and the absence of sunlight during the austral winter—severely compromises station maintenance and power supply continuity. As a result, the datasets are often discontinuous or limited in volume. A primary scientific challenge lies in the characterization of seismic events and the precise discrimination between icequakes (generated by ice fracturing and glacial dynamics) and signals of magmatic/hydrothermal/tectonic origin. In a closed-conduit system such as Mt. Melbourne, distinguishing deep-seated volcanic signatures from dominant glaciological noise is essential. Seismic precursors of potential unrest may be extremely subtle, infrequent, or masked by surface processes.

To overcome the limitations of manual classification and the scarcity of data, this study is currently building upon standard seismic analysis techniques to establish a solid diagnostic foundation. These classical methods are being used to identify and prepare the datasets that will, in the next phase, be processed through Machine Learning (ML) algorithms. By training these future models on specific spectral and temporal features extracted from the current analysis, the research aims to automate the characterization and cataloging of seismo-volcanic events. This evolving framework is designed to establish a reliable baseline of the volcano's background activity, providing the foundation for an automatic monitoring system capable of detecting subtle deviations and signalling potential volcanic unrest.

The ability to predict or promptly identify unrest at Mt. Melbourne is not only of scientific interest but it is of critical strategic importance. A significant unrest phase or eruptive event would profoundly impact nearby infrastructures, specifically the Italian Mario Zucchelli Station and the South Korean Jang Bogo Station.

Such an event could disrupt vital research activities, and pose a direct threat to the safety of personnel and the logistical networks that support international scientific cooperation in Antarctica. By establishing a robust monitoring framework, this research provides a vital tool for volcanic risk assessment and for the resilience of polar scientific operations.

Early detection of paroxysmal eruptions at Stromboli volcano using seismic data

Vittorio Minio[1,*], Flavio Cannavò[1], Marcello D'Agostino[1], Massimo Orazi[2], Eugenio Privitera[1]

[1]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Italy*
[2]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Napoli - Osservatorio Vesuviano, Italy*

**Corresponding author: vittorio.minio@ingv.it*

Paroxysmal eruptions at Stromboli represent a major hazard for both the volcano's summit area and the surrounding island. Early Detection (ED) systems are essential for mitigating these risks, but their effectiveness depends on delivering accurate and timely warnings, since false alarms or missed detections may have serious consequences for public safety and local infrastructure.

Previous geophysical studies have identified short-term precursory signals preceding paroxysmal eruptions at Stromboli, often occurring only minutes before the onset of the event. However, such limited lead times may be insufficient to ensure the safety of personnel working near the summit area. Improving risk mitigation, therefore, requires the identification of longer-term precursory signals associated with volcanic unrest.

Machine learning (ML) techniques offer a promising approach for the automatic recognition of such precursors in continuous geophysical data. ML models can detect complex patterns and temporal relationships that may precede eruptive events. By learning from historical observations, these models can support the real-time identification of anomalies and the early detection of impending explosive activity.

In this study, we explore the development of a new experimental ED system based on continuous seismic signals recorded by the permanent monitoring network of the Istituto Nazionale di Geofisica e Vulcanologia – Osservatorio Etneo during the 2017–2024 period. The proposed approach combines seismic observations with data-driven techniques to investigate potential mid- to long-term precursory patterns preceding paroxysmal activity at Stromboli.

The early recognition of such anomalies may provide a critical window for implementing preventive and evacuation measures. Beyond protecting lives and infrastructure, this type of system may improve emergency planning efficiency, reduce potential economic losses, and contribute to more effective volcanic risk management.

Clustering of tiltmeter data for automated identification of Mount Etna's active crater

Adriana Iozzia*, Marco Aloisi, Flavio Cannavò, Salvatore Gambino

INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Catania, Italy

**Corresponding author: adriana.iozzia@ingv.it*

Mount Etna (Italy) is an open-conduit basaltic volcano located on the eastern coast of Sicily and one of the most active systems globally, providing a natural laboratory for studying volcano dynamics. Its summit craters, Voragine (VOR), North-East Crater (NEC), Bocca Nuova (BN), and South-East Crater (SEC), usually display continuous degassing and intense Strombolian activity that sometimes culminates in powerful lava fountains. Since 2011, approximately 130 paroxysmal episodes have been documented, exhibiting a marked variability in the activation of different eruptive vents - from the frequent events at the SEC to the reactivations of the Voragine (2015, 2016, 2024), up to the rare NEC episodes recorded in December 2025. These events, due to the formation of dispersed ash plumes and fall-out deposits, may pose hazards to aviation and to population.

In volcano geophysics, continuous ground deformation monitoring is important for tracking magma ascent and system pressurization. Borehole tiltmeters, in particular, measure small changes in the ground's inclination (typically in the order of microradians). Mt. Etna's permanent tilt network includes 18 bi-axial instruments and a long-based tiltmeter, managed by Istituto Nazionale di Geofisica e Vulcanologia, Osservatorio Etneo. During lava fountains, these sensors recorded signal variations characterized by sharp changes (in the order of fractions of µrad, except for the ECPN station that showed larger amplitudes).

Traditionally, identifying the specific active crater relies on direct visual or thermal camera observations, which can be difficult under heavy cloud cover or adverse weather conditions. Furthermore, while seismo-acoustic networks (volcanic tremor and infrasound) provide real-time localizations, their spatial resolution can be constrained by the closeness of the summit vents, complex topography, and strong winds, sometimes introducing ambiguities in the localization of an accurate eruptive source. To overcome these limits, this study presents a framework integrating tiltmeter time-series with Machine Learning (ML) techniques.

We discuss the results obtained by applying both unsupervised and supervised ML algorithms to analyze and categorize the distinct deformation patterns generated by different craters.

By extracting specific features from the tilt signals - such as the maximum amplitude, the direction and azimuth of the deformation vectors, and the temporal waveform morphology - we evaluate the capability of clustering algorithms to group events without prior labeling, alongside the performance of supervised classification models trained to automatically identify the active vent.

Ultimately, integrating these ML models into operational monitoring routines could provide a reliable, weather-independent diagnostic tool, significantly enhancing near real-time source identification and volcanic hazard mitigation.

Seismology on rivers: a Machine Learning approach to monitor extreme floods generated by cyclonic storms

Mario Valerio Gangemi[1,*], Alfio Marco Borzì[1], Martina Allegra[2], Andrea Cannata[1,2], Flavio Cannavò[2], Concetto Spampinato[4], Luca Cavallaro[5], Enrico Foti[5], Rosaria Ester Musumeci[5], Stefano Parolai[3], Luca Zini[3], Matteo Picozzi[3], Salvatore Alparone[2], Danilo Contrafatto[2], Graziano Larocca[2], Francesco Panzera[1]

[1]*University of Catania, Department of Biological, Geological, and Environmental Sciences, Catania, Italy*
[2]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Catania, Italy*
[3]*University of Trieste, Department of Mathematics, Informatics and Geosciences, Trieste, Italy*
[4]*University of Catania, Department of Electrical, Electronic, and Computer Engineering, Catania, Italy*
[5]*University of Catania, Department of Civil Engineering and Architecture, Catania, Italy*

**Corresponding author: francesco.panzera@unict.it*

Over the last century, the frequency and intensity of extreme meteorological events have significantly increased, producing profound impacts on human societies worldwide. These events have resulted in substantial loss of life, extensive damage to infrastructure, and severe disruptions to transportation networks, agricultural systems, and natural ecosystems. Among the most hazardous phenomena, cyclonic storms—including tropical cyclones and Mediterranean hurricanes, commonly referred to as Medicanes—have become increasingly responsible for intense and persistent precipitation. Such conditions frequently trigger flood events, causing rapid and dangerous rises in river water levels and mobilizing large volumes of bedload during the most energetic episodes. The escalating occurrence of these extreme events has driven a growing scientific effort aimed at improving their characterization, monitoring, and prediction, with the ultimate goal of reducing associated risks and societal impacts.

Within this framework, the WAVEGUARD project focuses on enhancing monitoring capabilities and disaster response strategies to mitigate the effects of extreme meteo-marine events. Funded through the INTERREG VI-A Italia–Malta Programme and co-financed by the European Union, the project targets the areas surrounding the Maltese Islands and the channel between Malta and Sicily. By strengthening observational networks and data integration across this transboundary region, WAVEGUARD aims to support more efficient and timely responses to extreme meteorological and marine conditions, thereby improving regional resilience to natural hazards.

In recent years, artificial intelligence techniques have played an increasingly important role in advancing predictive models for extreme events. Machine learning and data-driven approaches have contributed to a deeper understanding of event dynamics and have supported the development of more effective early warning systems, even when faced with challenges such as noisy, sparse, or incomplete datasets.

In this study, we present a series of investigations that led to the development and characterization of a predictive model for river water height based on the integration of seismic and hydrological data. Our approach emphasizes the collection and analysis of data from rivers characterized by diverse geomorphological settings, hydrological regimes, and basin sizes. In addition, we deliberately employed seismic stations located at varying distances from the river channels, enabling a systematic comparison of different monitoring configurations. This strategy allows us to assess the influence of environmental conditions and sensor placement on signal quality and model performance, and to identify the most informative seismic features for water-level prediction.

The proposed system integrates machine learning algorithms designed to analyse multivariate time series, providing a foundation for near–real-time flood forecasting and risk assessment that relies primarily on seismic observations. By reducing dependence on in situ hydrometric sensors alone, this methodology offers a valuable alternative for regions that are particularly vulnerable to cyclonic storms and rapid water accumulation within drainage basins. As such, it represents a promising contribution to the development of resilient and adaptable monitoring systems for flood-prone environments.

Reconstructing significant wave height from microseism data using Machine Learning techniques: Sicily Channel case study

Alfio Marco Borzì[1,*], Vittorio Minio[1], Andrea Cannata[1,2], Francesco Panzera[1], Susanna Saitta[2], Salvatore Alparone[2], Giuseppe Ciraolo[3], Danilo Contrafatto[2], Sebastiano D'Amico[4], Giuseppe Di Grazia[2], Graziano Larocca[2], Flavio Cannavò[2]

[1]*University of Catania, Department of Biological, Geological, and Environmental Sciences, Catania, Italy*
[2]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Catania, Italy*
[3]*University of Palermo, Department of Engineering, Palermo, Italy*
[4]*University of Malta, Department of Geosciences, Msida, Malta*

**Corresponding author: alfiomarco.borzi@unict.it*

Microseism is the most continuous and ubiquitous seismic signal on the Earth, generated by the interaction between the sea waves and the solid Earth. Based on its source mechanism and spectral content, microseism can be divided into three main frequency bands, ranging from 0.05 to 0.4 Hz. Several authors used this seismic signal to retrieve information about sea state conditions and, specifically, about the significant wave heights.

In the framework of the WAVEGUARD project (Interreg Italia-Malta), we built a machine learning model able to predict the sea state in terms of wave height by using microseism. In particular, we implemented a supervised machine learning regression technique designed to analyse the non-linear relationship between seismic features and sea state parameters. The input dataset consisted of hourly root-mean-square (RMS) amplitudes computed in 14 frequency bands for each of the 14 broadband seismic stations used in this study, returning 588 predictive features. The target variable was the significant wave height available from the MEDSEA_MULTIYEAR_WAV_006_012 Copernicus hindcast products, also used to train our machine learning model.

Preprocessing steps were applied to ensure model robustness. In particular, features with excessive data gaps were removed, while limited missing values were interpolated using linear interpolation. Seismic signals contaminated by strong regional or teleseismic earthquakes were excluded using the USGS earthquake catalogue. Box–Cox transformations were applied to skewed variables, and all features were normalized to avoid bias linked with the differences in magnitude and measurement scale across variables. To avoid temporal leakage and overfitting, the four-year dataset (2018–2021) here analyzed was split into 40 non-consecutive temporal chunks, with 70% used for training and 30% for testing. A k-fold cross-validation (k=3) strategy was adopted during both hyperparameter tuning and model training.

In this work, we tested three regression algorithms: Random Forest (RF), Light Gradient Boosting (LGB), and K-Nearest Neighbors (KNN). Among them, Random Forest demonstrated the best predictive performance, likely due to its robustness to noise, ability to model non-linear relationships, and resistance to overfitting. The final RF model reached a coefficient of determination (R^2) of 0.89 and a mean absolute error of approximately 0.21 m during the validation period. While prediction accuracy was high under normal sea conditions, a slight underestimation was observed during extreme wave events, probably due to the limited number of extreme events within the training dataset.

These results demonstrate that microseism-based machine learning models can be considered as a complementary tool for near-real-time sea state monitoring and hence support the development of innovative seismic-driven marine monitoring systems to be integrated with the classic methodologies already implemented for this purpose.

Hidden simplicity behind seismic dispersion curves revealed by unsupervised learning

Han Zhang*, Jiao Yang, Sidao Ni, Zefeng Li, Xuewei Bao, Lijun Liu, Brandon Schmandt

Innovation Academy for Precision Measurement Science and Technology, China

**Corresponding author: hanzhang@apm.ac.cn*

This work explores the application of unsupervised deep learning, specifically Variational Autoencoders (VAEs), to extract fundamental features from high-dimensional seismological data. Traditionally, characterizing the Earth's lithosphere involves complex, multi-step inversions of surface-wave dispersion data that often suffer from overparameterization and subjective prior assumptions. Our research demonstrates that VAEs can objectively compress complex Rayleigh-wave phase velocity dispersion maps into a minimal set of two latent variables. This reveals an inherent coupling between crustal and mantle seismic properties. Our findings suggest that deep learning provides a robust, data-driven framework for objective tectonic regionalization and can streamline future multi-physics joint inversions by providing a more compact parameterization of the Earth's interior.

Physics-guided Machine Learning for meteotsunami detection from co-located microbarometer and tide-gauge records at Mazara del Vallo, Italy

Salvatore D'Amico[1,2,*], Salvatore Aronica[4,5], Andrea Cannata[2,3], Francesco Panzera[2]

[1]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione Roma 2, Rome, Italy*
[2]*University of Catania, Department of Biological, Geological, and Environmental Sciences, Catania, Italy*
[3]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Sezione di Catania - Osservatorio Etneo, Catania, Italy*
[4]*University of Palermo, Department of Engineering, Palermo, Italy*
[5]*CNR | Consiglio Nazionale delle Ricerche, Rome, Italy*

**Corresponding author: salvatore.damico2@ingv.it*

Meteotsunamis are long ocean waves of atmospheric origin in the "tsunami-like" band (minutes to a few hours), generated when propagating atmospheric disturbances transfer energy to the ocean. The coastal amplitude is then shaped by resonant and amplification mechanisms over the continental shelf and by the intrinsic response of bays and harbors. In the modern framework, these events are often viewed as the most impactful component of nonseismic sea-level oscillations at tsunami timescales (NSLOTT), with increasing attention to forcings capable of producing coherent responses over long distances [Vilibić et al., 2025; https://doi.org/10.1029/2024RG000867].

From an operational monitoring perspective, observational chains are now available that integrate minute-resolution tide gauges, microbarometers, satellite observations and numerical modeling. Multi-level schemes have been proposed to identify favorable synoptic conditions and support nowcasting [Vilibić et al., 2016; https://doi.org/10.3389/fmars.2016.00057], alongside high-resolution coupled regional forecasting approaches that demonstrate the feasibility of predicting selected events in specific settings [Makrygianni et al., 2025; https://doi.org/10.5194/egusphere-2025-3555]. Nevertheless, operational detection remains limited because the meteotsunami signal is often masked by tides/seiches and ordinary meteo-ocean variability, the coastal response is strongly site-specific and threshold-based criteria produce numerous false alarms.

In this context, within the WAVEGUARD project (Enhanced monitoring and disaster response to mitigate impacts of extreme meteo-marine events), funded by the Interreg Italy-Malta Program, we propose a study based on multi-year, co-located atmospheric pressure and sea-level records from Mazara del Vallo (Italy), where several meteotsunami time windows are already documented in the literature. Our goal is to develop a detector that outputs a ranked list of candidate meteotsunami events, assigning each a probability score and a confidence estimate, with the practical aim of reducing false positives relative to traditional threshold-based baselines in continuous monitoring. The approach is grounded in the analysis of the marine response to atmospheric forcing and explicitly includes estimates of coherence between the two signals. Specifically, we exploit: (i) sea-level indicators sensitive to rapid changes after removal of the low-frequency component dominated by the astronomical tide; (ii) energy measures in the meteotsunami band; and (iii) pressure–sea-level coupling metrics (cross-correlation and optimal lag) that robustly summarize event plausibility.

We further adopt an interpretable supervised classifier based on gradient-boosted decision trees (GBDT). Such models are well suited to imbalanced datasets and a limited number of true events (as in the case of meteotsunamis) and they can capture non-linear interactions between forcing and response (e.g., extreme marine events that are not meteotsunamis in

the absence of a coherent atmospheric forcing). The interpretability of GBDT also enables quantifying the contribution of coupling-related features and calibrating the probabilistic score for different alert levels, controlling the trade-off between detection and false alarms. Looking ahead, the detector can be integrated into multi-level monitoring chains and transferred to other sites/years through local recalibration, supporting comparative analyses of extreme NSLOTT processes.

DeepPinch: few-shot stratigraphic pinch-out detection in seismic images via synthetic pre-training

Shimon Cohen, Amotz Agnon*, Michael Werman

The Hebrew University of Jerusalem (HUJI), Jerusalem, Israel

Corresponding author: shimon.cohen3@mail.huji.ac.il

Automated interpretation of seismic reflection data is critical for subsurface geological analysis, yet the precise localization of stratigraphic features like pinch-outs remains an open challenge. The primary bottleneck is the prohibitive cost of generating accurate, pixel-level annotations, which traditionally requires laborious expert interpretation of full 3D seismic volumes. In this paper, we propose a few-shot domain adaptation framework to detect pinch-outs in 2D seismic images utilizing synthetic data. We introduce three distinct frameworks for synthetic data generation, enabling the creation of expansive, diverse, and precisely annotated datasets for pre-training. We utilize an Attention U-Net architecture equipped with Instance Normalization and Dropout to predict localized pinch-out probabilities. Our extensive experiments demonstrate that pre-training on a diverse synthetic corpus, followed by fine-tuning on a severely limited dataset of semi-supervised real seismic slices, yields significant performance gains. This two stage pipeline drastically outperforms models trained exclusively on real or synthetic data, successfully bridging the domain gap and providing a scalable solution for fine-grained stratigraphic feature extraction.

Crustal fracture classification and parameter estimation from InSAR displacement fields with interpretable descriptors

Louise Lizé[1,*], Christophe Lin-Kwong-Chon[1], Valérie Cayol[2]

[1]*Laboratoire d'Informatique, Systèmes, Traitement de l'Information et de la Connaissance (LISTIC), University Savoie Mont Blanc, Annecy-le-Vieux, 74944, France*

[2]*Université Clermont Auvergne, Laboratoire Magmas et Volcans (LMV), CNRS, IRD, OPGC, Clermont-Ferrand, France*

**Corresponding author: louise.lize@univ-smb.fr*

Satellite-based remote sensing enables large scale and precise monitoring of natural hazards. In particular, InSAR measures surface displacements that can be analyzed by inverse modelling to provide information on the processes driving deformation. Because inverse modelling is computationally costly, we investigate a machine learning approach to identify deformation source types from displacement fields, distinguishing magmatic intrusions from tectonic faults, while also exploring the estimation of physical parameters.

Inspired by the classification framework of Cayol et al. [2014], which uses ratios of surface displacement extrema to support mechanism diagnosis and parameter estimation, we introduce a machine learning based approach for more robust mechanism classification and parameter estimation. We compute interpretable descriptors from displacement fields, including global statistics and geometric characteristics of the displacement patterns, and train a random forest for classification. We also use random forest regression to explore the estimation of associated physical parameters.

The synthetic dataset used is generated from the Okada elastic model. In this model, the deformation source is a rectangular dislocation described by physical parameters. The resulting classes are well separable and the model achieves good performance, but considering simplified assumptions: a north-south fracture, a blind fracture and no atmospheric noise. The next step is to improve robustness under more realistic conditions. Perspectives include the test of others algorithms and feature selection strategies, before moving to real data and exploring deep learning approaches as complexity increases.

Using neural networks for InSAR images inverse modeling

Gregoire Vest*, Yajing Yan, Fabien Albino

Université Savoie Mont Blanc, Laboratoire d'Informatique, Systèmes, Traitement de l'Information et de la Connaissance (LISTIC), Annecy-le-Vieux, France

**Corresponding author: gregoire.vest@univ-smb.fr*

During volcanic unrest, inverse modeling of InSAR surface displacements is often used to retrieve the parameters of the source characteristics of volcanoes. Such approach could be costly in terms of computing time, and therefore not always adapted to routine modeling. Here, we are interested in exploring the potential and limitations of deep learning for the characterization of volcanic sources. We rely on our previous proof-of-concept work [Lopez Uroz et al., 2024], which has demonstrated the potential of neural networks in case of a point source model [Mogi, 1958].

This work will be used as a basis for expanding research into the benefits of deep learning in volcanic inverse modeling. In this study, we address issues of image resolution, patch dimension, and SAR acquisition geometry in order to properly prepare the learning dataset and to further improve the prediction accuracy of neural networks.

To do this, we continue with the Mogi model [Mogi, 1958] described by the depth of the volcanic source and its volume change. We propose multi-channel input, including multi-resolution and multi-geometry (i.e. descending and ascending) displacements maps issued from Interferometry Synthetic Aperture Radar (InSAR) in the case of previously deployed ResNet to estimate the depth and the volume change of the volcanic source. For results evaluation, besides the mean squared error of estimated parameters compared to the truth, we also consider structural similarity (SSIM) between the input displacement map and the reconstructed one as metric. Moreover, we perfome a sensitivity analysis in order to highlight the impact of resolution on each estimated parameters (i.e. volcanic source depth and volume change).

The multi-channel ResNet model is trained and first evaluated by synthetic datasets where the ground truth allows for quantitative assessment of the learning performance. The trained model is also applied to Sentinel-1 InSAR datasets covering the Suswa volcano during the period of 2016-2020 [Albino & Biggs, 2021]. The results are compared to those of the conventional inversion approach for testing.

According to the obtained results, the input image resolution to the neural network has significant impact on its effectiveness. As shown by the sensitivity analysis, the depth of the volcanic source is more sensitive to displacement information near the volcanic source, while its volume change is more sensitive to displacement information far from the volcanic source. Given the fixed input image size, a multi-resolution input therefore seems optimal to ensure both information. On the other hand, taking both SAR acquisition geometries, ascending and descending as input to the ResNet model is highly interesting for improving the inversion performance.

Template matching using Deep Learning

Dario Jozinović[1,*], Tania Toledo[1], Verena Simon[1], Toni Kraft[1], Frederick Massin[1], John Clinton[1], Olivier Meyers, Amilcar Roca[2]

[1]*Swiss Seismological Service (SED), ETH Zurich, Zurich, Switzerland*

[2]*Instituto Nacional de Sismología, Vulcanologia, Meteorologia e Hidrología (INSIVUMEH), Guatemala City, Guatemala*

**Corresponding author: dario.jozinovic@sed.ethz.ch*

The use of template matching is widespread in seismology (for e.g. earthquake detection). The most common technique to do template matching is cross-correlation - calculating Pearson correlation on a moving window. However, cross-correlation is a linear measure of similarity and can be sensitive to noise bursts and waveform distortions. Contrastive learning is a field of machine learning which aims to learn how to best transform inputs as to allow separation between similar and dissimilar data using a certain distance metric (e.g. Euclidean distance). The advantage of this learned similarity is that it can be optimized to measure non-linear relationships between data groups, give more importance to some parts of the template, etc. We develop a Siamese Neural Network (SNN) to improve template matching detections coming from a classical cross-correlation framework. More specifically, we use a dataset of manually checked true and false detections of this framework on several earthquake sequences in Switzerland. We then train our SNN to discriminate between these true and false detections and reduce the need for manual checks of the cross-correlation detections, showing improved performance over the classical cross-correlation techniques. In addition, we also show the capability of SNNs to perform in training-data-poor scenarios: measuring similarity and dissimilarity between the waveforms allows us to construct a large dataset of training pairs from small waveform datasets. To demonstrate the usefulness of SNNs in small dataset scenarios we show how an SNN can be used to detect lahars on volcano Santiaguito in Guatemala and provide early warning to the local population. To the same purpose, we also show an application of the method to detection of small rockfalls in Switzerland.

Unsupervised discrimination of earthquakes and blasts in Norway using representation learning

Miguel Neves*, Lars Ottemöller, Stéphane Rondenay

University of Bergen, Department of Earth Science, Bergen, Norway

**Corresponding author: miguel.neves@uib.no*

Accurately distinguishing natural earthquakes from anthropogenic seismic sources is essential for compiling reliable earthquake catalogs. This is particularly important in intraplate regions like Norway where low to moderate magnitude seismicity occurs alongside abundant anthropogenic sources, such as quarry blasts. Traditional supervised deep-learning approaches rely heavily on labeled training data, which can be limited or biased. To address this, we test a self-supervised strategy for discriminating earthquakes and blasts using seismic data from the Norwegian National Seismic Network.

Here we apply a representation-learning framework based on Bootstrap Your Own Latent (BYOL), which learns latent seismic features by aligning augmented views of the same waveform, eliminating the need for ground-truth labels. The model receives no prior information beyond the training hyperparameters. After training, we apply K-means clustering to the learned embeddings, providing only the number of expected clusters (two). Using these self-supervised representations, derived from 69,672 spectrograms corresponding to 20,894 events detected in Norway, we obtain a clear distinction between earthquake and blast signals, achieving an accuracy of 0.94 when evaluated against all available labels (confirmed and suspected), and 0.96 when limited to confirmed labels. As an additional test, we freeze the BYOL feature extractor and train a lightweight supervised classifier on the fixed embeddings using event labels. This simple supervised layer slightly increases accuracy to 0.97 on the confirmed dataset, showing that even a minimal classifier can effectively use the structure learned in the unsupervised stage.

Inspection of the embedded clusters reveals likely misclassified events, highlighting the potential of self-supervised learning to improve earthquake catalog quality even when labeled datasets are sparse or biased. Future work will analyze the learned latent features to gain insights into the physical attributes driving the separation between earthquakes and blasts.

Integrating game engine and PINN for multi-scale building cluster earthquake scenarios

Chenrui Fan*, Chen Gu, Zhuoyu Chen, Yuheng Wu, Zhihua Zhao, Yichen Zhong, Jiawei He, Peng Wu, Zhi Yuan, Yufeng He, Yifu Cheng

Tsinghua University, China

**Corresponding author: fcr22@mails.tsinghua.edu.cn*

Earthquakes pose significant risks to urban infrastructure and human safety, highlighting the need for advanced tools in disaster scenario reconstruction and risk mitigation. However, traditional earthquake simulation models often lack interactivity and dynamic adaptability, which limits their application in emergency response planning. Although rapidly evolving machine learning models show potential for real-time seismic damage prediction, their deployment is hindered by a scarcity of high-fidelity, large-scale datasets detailing the physical responses of complex urban building clusters.

To address this gap, this study develops and expands an interactive disaster scenario reconstruction system based on the Unity game engine. By integrating discrete wavenumber-based seismic simulations with game engine rendering and GPU acceleration technologies, the system creates a highly dynamic digital twin environment. In a case study spanning both a single building and a localized urban cluster, real-world geographic information and Building Information Modeling (BIM) data were imported. Furthermore, optical fibers were virtually deployed within the structures to capture Distributed Acoustic Sensing (DAS) data, facilitating the simulation and evaluation of a series of hypothetical earthquakes in the vicinity. However, calculating the complex dynamic responses of the building clusters creates a severe computational bottleneck, rendering traditional numerical methods impractical for real-time emergency applications.

To overcome this computational barrier, an end-to-end Physics-Informed Neural Network (PINN) was constructed and trained, leveraging the high-quality physical data generated by the engine. This deep learning architecture directly establishes a mapping relationship between ground motion model parameters and the dynamic responses of complex building clusters. Results indicate that this end-to-end surrogate model significantly reduces the computational overhead of traditional dynamic analysis methods while maintaining the consistency of prediction outcomes. Ultimately, this project enables the real-time visualization of complex disaster scenarios, offering decision-makers a novel solution for post-earthquake emergency response and intelligent disaster prevention and mitigation.

Diffusion-based denoising for submarine DAS seismic monitoring

Jiawei He*, Chen Gu, Yichen Zhong, Peng Wu, Zhi Yuan, Chenrui Fan, Shigui Zhang, Jiachao Xu, Xiaoyong Qin, Xiaoyan Fan, Hongbo Sun

Tsinghua University, Department of Civil Engineering, China

**Corresponding author: hejw21@mails.tsinghua.edu.cn*

Submarine seismic monitoring is often limited by strong ambient and instrumental noise, which can obscure weak seismic signals and reduce the performance of automatic event detection. In this study, we investigate a diffusion-based denoising framework for Distributed Acoustic Sensing (DAS) data acquired from a submarine fiber-optic cable in the South China Sea. The DAS system converted the existing seafloor communication fiber into a dense sensing array and continuously recorded high-resolution spatiotemporal strain-rate signals over a two-month period, providing a valuable dataset for offshore seismic monitoring.

To improve the quality of the submarine DAS recordings, we developed a denoising method based on a diffusion model, which progressively reconstructs clean seismic waveforms from noisy observations through an iterative reverse process. The model was designed to suppress strong spatially correlated marine noise while preserving coherent seismic arrivals and weak event-related signals. The denoised results show clearer waveform continuity, enhanced signal coherence across channels.

To assess the practical impact of denoising, we further applied an automatic seismic event detection workflow to both the original and denoised datasets. The results show that the denoised data enable the detection of more seismic events than the raw data alone. In particular, weak events previously masked by background noise became identifiable after denoising, and the detected events exhibited more coherent multi-channel signatures, improving the reliability of event recognition.

These results demonstrate that diffusion-based denoising is an effective approach for enhancing submarine DAS data quality and improving downstream seismic event detection. This study highlights the potential of combining continuous offshore DAS observations with advanced generative denoising models to strengthen seismic monitoring capability in complex marine environments.

From faults to volumes: the development of an invertible Deep Learning surrogate for versatile kinematic source modelling

Jonathan Bedford*, Kaan Çökerim

Ruhr University Bochum, Institute of Geosciences, Tectonic Geodesy, Bochum, Germany

**Corresponding author: jonathan.bedford@rub.de*

For some of the world's most tectonically and volcanically active regions, advances in geodetic monitoring infrastructure and processing are building towards a future data scenario in which we can track the full 3D motion of the land surface at sub-kilometre resolution and sub-weekly timescales. Naturally, this presents a challenge to the community: relating the observed surface displacements to their underlying deformation sources. Such sources may be localised due to slip along faults or more distributed due to processes such as flow in less viscous pockets of the lithosphere.

Currently, deformation source models are typically constructed using prior structural geological information—for instance, positioning a fault source close to planes indicated by seismic catalogues and focal mechanisms, or extrapolating geologically mapped structures to depth. Optimisation of model parameters then occurs via direct inversion routines or through iterative adjustment of parameters when forward simulation models must be run. For example, post-seismic deformation following large earthquakes is often modelled with a coseismic "kick" and a subsequent numerical solution using Finite Element Modelling (FEM). In such "kick-and-flow" models, remeshing geometry between model runs is challenging, and the 3D solutions are generally too computationally expensive to allow exhaustive model parameter optimisation. In recent years, a solution to this problem was demonstrated using anelastic eigenstrain volumetric sources for off-fault deformation and traditional elastic half-space dislocation sources for fault motions, all combined in a linear system of equations.

In this presentation, we show how the volumetric deformation source can generalise, when shrunk to a plane, to represent fault motions. Given this convenient property of the volumetric source, we trained a deep learning surrogate to reproduce surface deformation patterns from a variety of "fat", volumetric-like and "thin", dislocation-like source samples. The trained surrogate is invertible in the sense that we can recover source parameters from deformation maps by freezing the weights of the surrogate, adding a scale parameter, and solving for the input parameters to the network. Furthermore, we demonstrate the model's performance for synthetic cases using a sparse modelling approach that aims to fit the signal with as few deformation sources as possible. Generally, the surrogate is able to correctly invert the known synthetic source parameters, while the speed of the model in forward mode allows an efficient search of the solution space from which we can select advantageous starting points for the inversion.

In addition to presenting the capabilities of the versatile deformation source model, we will also share the lessons learned during the evolution of the surrogate training approach. A notable example concerns the treatment of sources that intersect the free surface: the original version of the model, which mimicked only elastic dislocation sources, handled such boundary-crossing sources differently from the later volume-based models. We hypothesise that the manner in which the model treats sources approaching the surface boundary influences the rate at which inversions converge to the true solution, as appropriate boundary handling can reduce solution non-uniqueness.

Blind surrogate inversion of deformation sources

Kaan Çökerim*, Jonathan Bedford

Tectonic Geodesy Working Group, Institute of Geoscience, Ruhr University Bochum, Germany

**Corresponding author: kaan.coekerim@ruhr-uni-bochum.de*

Geodetic surface displacement data from GNSS and InSAR provide valuable constraints on sub-surface deformation processes, but inferring the number, geometry, and deformation mode of the underlying sources remains a challenging inverse problem. This is especially true in multi-source settings, where repeated evaluations of forward models make blind exploration of the model space computationally expensive. We present a differentiable deep learning surrogate inversion framework based on learned cuboid anelastic deformation sources and sparse optimization. Trained on analytical cuboid half-space solutions, the surrogate provides a flexible source representation capable of approximating a range of deformation styles, from volumetric sources to near-planar fault-like geometries. Once trained, the surrogate serves as a fast differentiable forward model, enabling gradient-based optimization of source geometry and eigenstrain directly against observed deformation data. In contrast to conventional kinematic inversions that require source types and source numbers to be prescribed a priori, the deep learning surrogate supports blind multi-source recovery under sparsity-promoting constraints. Adding to earlier proof-of-concept results, we extend the surrogate inversion method to synthetic line-of-sight data and apply it to real InSAR observations. The results demonstrate the promise of physics-guided surrogate models as scalable tools for geophysical inverse problems and for data-driven interpretation of satellite geodetic measurements.

Source localization from submarine DAS data using deep Bayesian beamforming

HE Yufeng*, GU Chen, HE Jiawei, CHEN Zhuoyu, ZHONG Yichen, WU Peng, YUAN Zhi, FAN Chenrui, WU Yuheng

Tsinghua University, China

**Corresponding author: ml4seg@mail.sphenhe.me*

This study developed a U-Net-assisted Bayesian beamforming framework for source localization from submarine DAS data. Conventional beamforming often estimates source location by selecting the peak response on a scanning grid, but its performance can be sensitive to the choice of time window, frequency band, and phase content. It can also become less reliable when channel quality is uneven, noise levels vary across the cable, or multiple candidate peaks appear in the beamforming response. In addition, repeated scanning and parameter tuning may lead to high computational cost and limited interpretability. To address these problems, this study combined data-driven learning with probabilistic beamforming, so that the final output was not only a single location estimate but a probability distribution over candidate source locations, allowing the localization result to be expressed in a more flexible and informative way.

The proposed method followed a coarse-to-fine workflow. First, earthquake event windows were selected manually, and the corresponding submarine DAS data segments were extracted for analysis. A U-Net model was then used to learn two event-dependent quantities from the observed channel-time data: a reliability weight for each channel or subarray, and a prior distribution over candidate source locations on a predefined (r,θ) grid, where r denotes distance and θ denotes direction relative to the cable. The learned weights reduced the influence of noisy or low-quality channels, while the learned prior guided the search toward more likely source regions and reduced unnecessary scanning over the full grid.

Based on these U-Net outputs, a Bayesian beamforming stage was constructed. For each candidate location on the predefined (r,θ) grid, the DAS wavefield was aligned according to the corresponding propagation delay pattern, and a weighted beamforming likelihood was calculated. This likelihood was then combined with the U-Net-generated prior to produce a posterior probability distribution over source location. In this way, the U-Net did not replace beamforming, but improved it by providing adaptive weights and informative priors. The method therefore produced both the most likely source location and an associated uncertainty estimate, rather than relying only on a peak-based point prediction.

The results showed that this framework provided a practical and interpretable way to improve source localization from submarine DAS data. Compared with conventional equal-weight beamforming, it reduced sensitivity to noisy channels and parameter choices, narrowed the effective search region, and produced more stable localization results together with uncertainty information. These results suggest that combining U-Net-based feature learning with Bayesian beamforming is a useful approach for more robust and efficient submarine DAS source localization.

Earthquake seismogram denoising across time, time–frequency, and hybrid domain approaches

Niko Dahmen

Institute of Geophysics, ETH Zurich, Zurich, Switzerland

Corresponding author: nikolaj.dahmen@eaps.ethz.ch

Deep learning-based denoising models have shown great potential for separating earthquake and noise signals by processing data in either the time-domain or time-frequency domain. Here we present systematic benchmarks comparing models operating in these two domains alongside hybrid models - introduced here - that jointly leverage high-resolution waveforms and time-frequency decomposition. Models are trained on a dataset recorded in Switzerland, curated specifically for the denoising tasks.

We evaluate performance across metrics and tasks critical for seismic monitoring: waveform quality, signal-to-noise ratio (SNR) improvements, signal detection, phase picking, and first-motion polarity determination. We further assess generalization on out-of-distribution data recorded in Italy and compare results against published denoising models.

Our results show that: (1) the time-domain model best capture waveforms at moderate-to-high SNRs, and naturally preserve high-resolution features critical for seismic data analysis, but degrade under low-SNR conditions; (2) time-frequency models improve detection at low SNRs through signal decomposition, but introduce windowing artifacts that reduce waveform quality and require additional processing; (3) the proposed hybrid model combine these advantages, accurately reconstructing signals across a broad SNR range and exhibiting the best overall performance in phase picking and first-motion polarity determination on both the benchmark and generalization datasets.

These results highlight that data representation is a consequential design choice in deep learning for seismology, alongside training data composition and network architecture. We offer practical guidance on model selection, discuss remaining limitations, and identify where deep learning denoising provides advantages beyond conventional processing.

This work is carried out within TRANSFORM², funded by the European Commission under project number 101188365 within the HORIZON-INFRA-2024-DEV-01-01 call.

Time delay neural networks reveal pressure-independent fault rupture processes in laboratory acoustic emission

Sergio Vinciguerra[1,2,*], Thomas King[3]

[1]*University of Turin, Department of Earth Sciences, Turin, Italy*
[2]*INGV | Istituto Nazionale di Geofisica e Vulcanologia, Rome, Italy*
[3]*University of Birmingham, Birmingham, UK*

**Corresponding author: sergiocarmelo.vinciguerra@unito.it*

Fault nucleation and growth are central to earthquake hazard. Here we analyse fault development in Alzo granite under triaxial confining pressures of 5–40 MPa using a time-delay neural network applied to multi-parameter acoustic-emission data. The model integrates waveform attributes (peak delay, scattering) with occurrence metrics (event rates, Gutenberg–Richter b-value, spatial fractal dimension) to track the transition from distributed microcracking to localised faulting. Genetic algorithms optimise the network, which dynamically weights parameters to characterise fault growth. We find three phases consistent across pressures: microcrack nucleation marked by scattering changes; fault growth captured by evolving spatial and magnitude distributions; and coalescence with rapid peak delay increases and b-value change. The model predicts stress-drop timing and size across pressures and failure mechanisms, from axial splitting to shear localisation, linking waveform features to physically interpretable phases of deformation.

N-dimensional field reconstruction from sparse observations

Javier E. Santos

Los Alamos National Laboratory, Los Alamos, NM, USA

Corresponding author: jesantos@lanl.gov

The reconstruction of complex time-evolving fields from sensor observations is a grand challenge. Frequently, sensors have extremely sparse coverage and limited onboard compute for measuring highly nonlinear phenomena. While numerical simulations can model some of these phenomena using partial differential equations, the reconstruction problem is ill-posed. Data-driven strategies provide crucial disambiguation, but these suffer in cases with small amounts of data, and struggle to handle large domains. In this presentation, we will showcase our most recent advances in data-driven frameworks across a range of applications, including subsurface property prediction (temperature, porosity, strain), tsunami wave reconstruction, locating lost wells using drones, optimal sensor placement, among others.

Autonomous research agents and structured feature operators for earthquake forecasting

Allie Hutchison*, Ryan Parks, Daniel Friedman, Adam Goldstein

Hum Lab, USA

**Corresponding author: allie@humlab.com*

Earthquake forecasting has been a north star for seismology for decades, yet operational results have remained elusive. Recent advances in AI have provided a jumpstart to the field, including automated research agents that accelerate model development, and mixture-of-experts architectures that combine the strengths of multiple modeling approaches. Leveraging these techniques, in four months, we developed a novel earthquake forecasting system that passes all five CSEP consistency tests (N, L, S, M, CL) for California earthquakes of magnitude 3.95 and greater.

The AI-accelerated development workflow is a central contribution of this work. Autonomous research agents handled code generation, experimental iteration, and evaluation pipelines, enabling rapid exploration of the model design space at a pace that would have been impractical through traditional development alone. However, responsible use of these tools demands careful attention to validation rigor, scientific reproducibility, and the risk of automation bias, where the speed and fluency of AI-generated outputs can outpace human verification. We discuss the guardrails our team adopted, including mandatory walk-forward testing, human-in-the-loop review of all model decisions, and explicit separation of AI-generated code from scientific interpretation. We believe this human–AI collaborative workflow offers a template for research teams seeking to tackle problems previously requiring far greater resources.

The forecasting system integrates multiple data modalities (including earthquake catalogs, fault geometry, seismic data, stress-related variables, and geodetic observations) within a structured representation framework. This framework maps heterogeneous inputs into a unified latent representation, enabling consistent feature construction across data sources without reliance on extensive manual tuning across modalities. Structured feature operators extract informative signals for downstream forecasting within a machine learning pipeline evaluated using daily walk-forward testing, with graph neural networks and graph transformers on the horizon. The representation is designed to balance expressivity and computational tractability, allowing efficient integration with machine learning workflows.

Adaptive Deep Learning approach for laboratory earthquake forecasting using multi-channel acoustic emission data

Ke Gao*, Zhixiong Leng

Southern University of Science and Technology, Shenzhen, Guangdong, China

**Corresponding author: gaok@sustech.edu.cn*

Multi-channel acoustic emission data, typically collected during laboratory fault shear experiments, play a crucial role in analyzing fault instability and provide valuable insights into the physical processes preceding failure. Nevertheless, the spatial and temporal migration of nucleation and slip event failures during fault evolution results in the dynamic redistribution of critical information among multiple acoustic emission channels, producing significant non-stationarity in observational signals. Additionally, sensor anomalies and noise contamination further complicate channel observations, thereby increasing modelling complexity and presenting substantial challenges for conventional static modelling methods. To address these issues, this study introduces an adaptive multi-channel deep learning approach for forecasting shear stress in laboratory faults. The proposed framework incorporates shared-weight feature extraction and adaptive channel weighting, allowing the model to learn consistent cross-channel representations while dynamically prioritizing informative and reliable channels. This design enhances the model's ability to adapt to the migration of essential information across acoustic emission channels and improves robustness in feature learning amid sensor disturbances. Consequently, non-stationary precursor information is adaptively represented, enabling effective characterization of shear stress evolution. Validation across diverse laboratory datasets indicates that the proposed method surpasses traditional models in both shear stress prediction accuracy and model stability, highlighting its efficacy in capturing shear stress evolution under complex observational scenarios.

GPU accelerated SEAS simulations for Machine Learning applications

Gabriele Benedetti*, Elías Rafn Heimisson

University of Iceland, Iceland

**Corresponding author: gab28@hi.is*

Sequences of Earthquakes and Aseismic Slip (SEAS) simulations focus on km scale models and consider all phases of faulting from aseismic slip to earthquake nucleation, propagation and termination. Many codes exist that use different approaches to tackle SEAS simulations; however, solutions designed to leverage the potential of GPUs to parallelize and speed up the simulation are limited. While these simulations often produce simple and periodic earthquake cycles, recent efforts show that at very large spatial scales, the complexity of the simulations starts to mimic that of natural faults. We use a GPU parallelized SEAS quasi-dynamic solver written in Julia (HighSEAS.jl, Benedetti and Heimisson [2026]) based on the Spectral Boundary Integral Method (SBIM) to efficiently run high resolution simulations at extreme length and time scales, thus giving the possibility to simulate fault ruptures resolved at realistic scales with high L/Linf ratios. This improved efficiency and expanded parameter range make SEAS a powerful tool for generating controlled synthetic and physically realistic earthquake catalogs that can be used to train machine learning algorithms. Systematically tuning the L/Linf ratio offers a controlled way to regulate the simulations complexity and thus offering various levels of challenges for ML prediction. We propose to perform a series of simulations with increasingly higher L/Linf ratios and use this output to train a simple neural network. We then evaluate the predictive capabilities of the increasingly complex data and identify the most valuable data components that can be used to forecast the simulated full fault ruptures, i.e. the mainshocks of the system.

Meta-learning for few-shot forecasting of laboratory earthquakes

Zhixiong Leng*, Ke Gao

Southern University of Science and Technology, Department of Earth and Space Sciences, Shenzhen, Guangdong, China

Corresponding author: lengzx2024@mail.sustech.edu.cn

Laboratory earthquake prediction commonly utilizes frictional stick-slip experiments to gather observational data on fault development. In contrast to the repeatable and controllable data acquisition afforded by laboratory environments, natural earthquakes are often limited by short observation periods and a lack of high-quality precursor records. Consequently, the number of effective samples for model training is extremely low, making it challenging to fully train complex deep learning models and restricting the application of data-driven approaches to real-world fault systems. Thus, effectively leveraging controlled laboratory data to create predictive models with cross-condition and cross-scale transfer capabilities in few-shot scenarios has become an essential scientific issue connecting laboratory and natural earthquake prediction. To address this challenge, the present study integrates a meta-learning approach within a deep learning prediction framework, proposing a few-shot adaptive learning model for laboratory earthquake forecasting. This method constructs a shared deep feature representation across tasks, enabling the model to identify consistent patterns of fault evolution from multiple frictional stick-slip experiments. Utilizing task-level rapid parameter adaptation inherent to meta-learning, the model achieves rapid convergence and accurate predictions even with limited sample sizes under new experimental conditions. Comparative evaluation on diverse datasets with heterogeneous experimental settings demonstrates that the proposed method greatly enhances prediction accuracy and generalization performance in few-shot settings, while offering improved stability in cross-condition transfer tasks. These results provide important methodological support for advancing laboratory earthquake prediction models toward application in natural earthquake scenarios.

Automatic analysis of seismic waveforms for sea-ice properties estimation

Jonas Michael*, Ludovic Moreau, Marielle Malfante

Université Grenoble Alpes, ISTerre, CEA, France

**Corresponding author: jonas.michael@univ-grenoble-alpes.fr*

The estimation of sea-ice properties is a scientific issue of increasing significance for better monitoring, understanding and predicting the dynamic of a warming Arctic ocean.
Among the various monitoring tools, Bayesian approaches analyzing seismic waveforms collected on the field have allowed to estimate sea-ice thickness and its mechanical properties with high temporal and spatial resolution. However, such analyses are costly, preventing in situ application and real-time monitoring.
To reduce the cost of this inversion problem we consider several scenarios of increasing interest for field deployment to train frugal neural approaches.
Convolutional neural networks are trained via supervised learning on synthetic datasets of varying sizes, generated using a physically informed forward model that propagates input pulses across diverse ice thicknesses, source-receiver distances and array sizes. We evaluate varying model architectures, dataset scales, and required receiver counts, notably finding that frequency-domain representations of the input waveforms improve network performance compared to time-domain inputs.
Additionally, the developed networks are validated across distinct test scenarios and confronted with raw field data. While demonstrating high accuracy on physics-informed synthetic data, predictive performance degrades significantly on real-world waveforms. Further investigation reveals this discrepancy stems from a pronounced gap between the propagation model and field data rather than unfamiliar source characteristics. The neural networks become tightly fitted to the physical equations used to generate the training data, demonstrating that these synthetic models, although accurately reproducing field data, fail to generate training data that capture the complexity of seismic wave propagation.

QUADERNI di GEOFISICA

ISSN 1590-2595 https://collane-editoriali.ingv.it/index.php/qdg/it/index/

QUADERNI di GEOFISICA è una rivista scientifica che pubblica dati, misurazioni e osservazioni geofisiche (anche preliminari) in italiano e inglese. L'obiettivo è una diffusione rapida e gratuita per la comunità scientifica, garantita da un processo di peer-review e da un Editorial Board multidisciplinare. La rivista è inclusa in "Emerging Sources Citation Index" di Clarivate Analytics e "Open Access Journals" di Scopus, e adotta il modello Diamond Open Access, che non prevede costi né per gli autori né per i lettori.

QUADERNI di GEOFISICA is a scientific journal that publishes geophysical data, measurements and observations, including preliminary results, in Italian and English. Its aim is to ensure rapid and free dissemination to the research community, supported by a peer-review process and by a multidisciplinary Editorial Board. The journal is included in the Emerging Sources Citation Index of Clarivate Analytics and it's included in the Open Access journals of Scopus. It follows the Diamond Open Access model, with no charges for either authors or readers.

RAPPORTI TECNICI INGV

ISSN 2039-7941 https://collane-editoriali.ingv.it/index.php/rt/it/index

RAPPORTI TECNICI INGV è una rivista tecnico-scientifica che mette in mostra le numerose attività tecnologiche e di ricerca condotte dai ricercatori e tecnici INGV, e non solo. Pubblicata in italiano e inglese, RAPPORTI TECNICI INGV è ideata per valorizzare anche quei lavori che offrono un contributo significativo dal punto di vista tecnico, operativo o documentale fondamentale per lo sviluppo dell'Ente. La rivista adotta il modello Diamond Open Access, che non prevede costi né per gli autori né per i lettori.

RAPPORTI TECNICI INGV is a technical–scientific journal that showcases the wide range of technological and research activities carried out by INGV researchers and technical staff, as well as by external contributors. Published in Italian and English, RAPPORTI TECNICI INGV is designed to highlight works that provide a significant contribution from a technical, operational, or documentary perspective, and that are essential to INGV development. The journal follows the Diamond Open Access model, with no charges for either authors or readers.

MISCELLANEA INGV

ISSN 2039-6651 https://collane-editoriali.ingv.it/index.php/mi/it/index

MISCELLANEA INGV rappresenta una tipologia di editoria non tradizionalmente catalogabile nelle ordinarie riviste scientifiche di settore. Pubblicata in lingua italiana e inglese, MISCELLANEA INGV offre uno spazio per la pubblicazione di materiali eterogenei che, pur non rientrando nei formati più strutturati delle altre Collane editoriali INGV, rivestono un interesse scientifico, tecnico o divulgativo. La rivista adotta il modello Diamond Open Access, che non prevede costi né per gli autori né per i lettori.

MISCELLANEA INGV represents a form of publishing that cannot be classified within the standard scientific journals. Published in Italian and English, MISCELLANEA INGV provides a venue for the publication of heterogeneous materials that are of scientific, technical or outreach interest, even though they do not fall within the more structured formats of the other INGV editorial series. The journal follows the Diamond Open Access model, with no charges for either authors or readers.

Progetto grafico
Barbara ANGIONI
Istituto Nazionale di Geofisica e Vulcanologia

Ufficio Editoriale
Francesca DI STEFANO | Coordinatore

Segreteria di redazione
segreteria.collane-editoriali@ingv.it

Editing | Produzione | Grafica redazionale
Barbara ANGIONI
Massimiliano CASCONE
Francesca DI STEFANO
Patrizia PANTANI

© 2026 the Author(s). Open Access.
This article is licensed under a Creative Commons Attribution 4.0 International License.

Creative Commons Attribution 4.0 International

ML4SEG



ISTITUTO NAZIONALE **DI GEOFISICA E VULCANOLOGIA**