

# Seismic detection using submarine cable polarization signals with machine learning

---

Received: 20 August 2025

Accepted: 12 March 2026

Cite this article as: Caruso, M., Morelli, M., Monaco, A. *et al.* Seismic detection using submarine cable polarization signals with machine learning. *Commun Earth Environ* (2026). <https://doi.org/10.1038/s43247-026-03434-x>

Mario Caruso, Michele Morelli, Alfonso Monaco, Nicola Amoroso, Luigi Avagliano, Tommaso Maggipinto, Danilo Decaroli, Alberto Marullo, Pietro Patimisco, Angelo Sampaolo, Francesco Simeone, Vincenzo Spagnolo, Andrea Zifarelli, Loredana Bellantuono & Roberto Bellotti

---

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

# Seismic detection using submarine cable polarization signals with machine learning

Mario Caruso<sup>1,+</sup>, Michele Morelli<sup>1,+</sup>, Alfonso Monaco<sup>\*1,2</sup>, Nicola Amoroso<sup>4,2</sup>, Luigi Avagliano<sup>8</sup>, Tommaso Maggipinto<sup>1,2</sup>, Danilo Decaroli<sup>5</sup>, Alberto Marullo<sup>5</sup>, Pietro Patimisco<sup>6</sup>, Angelo Sampaolo<sup>6</sup>, Francesco Simeone<sup>7</sup>, Vincenzo Spagnolo<sup>6</sup>, Andrea Zifarelli<sup>6</sup>, Loredana Bellantuono<sup>3,2,++</sup>, and Roberto Bellotti<sup>1,2,++</sup>

<sup>1</sup>Università degli Studi di Bari Aldo Moro, Dipartimento Interateneo di Fisica “M. Merlin”, Bari, 70125, Italy

<sup>2</sup>Istituto Nazionale di Fisica Nucleare, Sezione di Bari, Bari, 70125, Italy

<sup>3</sup>Università degli Studi di Bari Aldo Moro, Dipartimento di Biomedicina Traslazionale e Neuroscienze, Bari, 70121, Italy

<sup>4</sup>Università degli Studi di Bari Aldo Moro, Dipartimento di Farmacia - Scienze del Farmaco, Bari, 70126, Italy

<sup>5</sup>Telecom Italia Sparkle S.p.A., Roma, 00125, Italy

<sup>6</sup>PolySense Lab, Dipartimento Interateneo di Fisica, University and Polytechnic of Bari, Via Amendola 173, Bari, 70126, Italy

<sup>7</sup>National Institute of Geophysics and Volcanology, INGV, Rome, Italy

<sup>8</sup>Marina Militare Italiana, Palazzo Marina in Piazza della Marina 4, Roma, 00196, Italy

\*alfonso.monaco@uniba.it

+These authors contributed equally to this work.

++These authors jointly supervised this work.

## ABSTRACT

Submarine fiber-optic communication cables are emerging as potential platforms for large-scale geophysical sensing. This study evaluates the feasibility of exploiting the state of polarization of optical signals in the live Med-Nautilus submarine cable for earthquake detection. We analyze state of polarization time series collected between 2022 and 2024 to assess whether seismic events induce detectable perturbations. Supervised and unsupervised Machine Learning models, including Logistic Regression, Extreme Gradient Boosting, and Deep Autoencoders, are applied for classification and anomaly detection. Results indicate that moderate-to-large earthquakes (magnitude  $\geq 5$ ) generate measurable state of polarization anomalies, although the cable was not designed for sensing. The Extreme Gradient Boosting model achieves up to 60% accuracy in distinguishing seismic signals from background noise. These findings provide the first empirical validation of state of polarization-based seismic detection under real operating conditions and support the potential of existing submarine cable networks as scalable, cost-effective geophysical monitoring systems.

## Introduction

Earthquake detection and early warning systems are essential tools to reduce seismic risk, particularly in tectonically active and densely populated coastal regions<sup>1–6</sup>. Although terrestrial seismic networks provide robust coverage on land, their reach across the oceans remains limited, despite the fact that a big portion of global seismic activity occurs beneath the seafloor. This observational gap is especially relevant in the Mediterranean Sea, where the convergence of complex tectonic boundaries and high population density demands enhanced offshore monitoring capabilities. In recent years, submarine fiber-optic communication cables, originally deployed for global data transmission, emerged as a promising medium for ocean floor seismic sensing. With a global reach exceeding  $1.3 \cdot 10^6$  km, these cables often cross tectonically active zones, offering a strategic asset for geophysical observation. Through techniques such as Distributed Acoustic Sensing (DAS) and optical polarization sensing<sup>7–13</sup>, it is possible to detect seismic signals by analyzing changes in back-scattered light or in the State of Polarization (SOP) of the transmitted optical signal<sup>14–16</sup>. Slader et al.<sup>17</sup> demonstrated that a DAS interrogator using a standard submarine telecommunication cable system can detect not only local and teleseismic events but also complex solid Earth–ocean interactions. However, at the current stage of development, this technique is limited to sensing cable segments shorter than 100 km and requires a dark fiber dedicated exclusively to measurement. Consequently, it is not yet feasible to apply this method to active transoceanic telecommunication cables, which typically span distances of thousands of kilometers. Complementing this result, Zhan et al.<sup>18</sup> provided compelling evidence that SOP-based techniques can detect both seismic waves and ocean

surface gravity waves over intercontinental distances. These findings underscore the transformative potential of leveraging existing transoceanic telecommunications infrastructure as a global-scale geophysical observatory, enabling continuous and cost-effective monitoring of submarine environments.

Building on these foundational studies, our research focuses on the application of SOP signals from Med-Nautilus, a submarine internet cable in the Mediterranean Sea, for earthquake detection. We analyzed the signals that an Italian telecommunication company, known as Telecom Italia Sparkle S.p.A., has provided from its Med-Nautilus submarine cable infrastructure between 2022 and 2024. Med-Nautilus is an advanced submarine cable network that connects key locations in the Mediterranean Sea, such as Italy, Greece, Turkey, Cyprus and Israel, with a high-performance fiber optic infrastructure, enabling ultra-fast connections. The network section that we analyzed has endpoints located in Catania (Sicily, Italy) and Tel Aviv (Israel). Data employed in this work were recorded in Catania. A key innovation of our approach is the integration of Machine Learning (ML) and Deep Learning (DL) methods to enhance the sensitivity, automation and real-time potential of signal analysis. Specifically, to classify signals and identify anomalies indicative of seismic activity, we employed Logistic Regression and Extreme Gradient Boosting (XGBoost), with the latter providing better performance, and Autoencoders, a form of unsupervised neural network. Although the use of ML in terrestrial seismology has seen increasing success<sup>19,20</sup>, its application to optical signal data from submarine cables remains largely unexplored.

In this work, we propose an investigative pipeline for the classification and detection of seismic events in the Mediterranean region. The analysis started with the preprocessing of time series describing the SOP of the electromagnetic signal from the Med-Nautilus submarine optical fiber cable. We first explored two traditional benchmark techniques for signal analysis and event detection. However, these approaches did not provide reliable results, thus motivating the transition toward data-driven ML and DL methodologies. We obtained promising results in the supervised classification of seismic events with magnitude  $M \geq 5$ . Additionally, we achieved partial success in the anomaly detection of seismic events characterized by the same magnitude scale. To the best of our knowledge, this is the first study that systematically investigates the SOP data of the Med-Nautilus submarine cable in the Mediterranean region through the application of ML and DL techniques to classify and detect seismic events, respectively.

The research described in this work was co-funded by the Italian Ministry of Defense under the National Military Research Plan (PNRM).

## Results

This section, structured into three parts reflecting the different objectives of the study, shows the results of the analyses carried out on time series describing the temporal evolution of the SOP along the Med-Nautilus optical fiber cable. The dataset analyzed in our work includes Stokes parameters and the entries of Jones matrix, both appropriately preprocessed. Figure 1 shows a simplified flowchart of the steps presented in this study.

The first subsection (*Benchmark methods analysis*) reports the results of two benchmark methods: STA/LTA (Short Time Average - Long Time Average) and fixed-threshold spectral energy analyses. This subsection explains the simplest approach employed to attempt the detection of the seismic events under investigation, relying solely on basic amplitude-and energy-based criteria without any learning or adaptive modeling.

The second subsection (*Machine Learning analysis*) shows the findings obtained from the ML analysis, which focuses on classifying seismic events; the approach relies on physical and statistical features extracted from SOP time series.

The third subsection (*Deep Learning analysis*) reports the results of the DL analysis, aimed at detecting seismic events treated as anomalies in the polarization state. Here, the analysis is based on time series of two independent Stokes components,  $S_1$  and  $S_2$ .

### Benchmark methods: STA/LTA analysis

We first assessed the feasibility of detecting seismic-induced perturbations using a classical STA/LTA detector applied to both Stokes parameters and Jones matrix components. Several configurations of short- and long-term averaging windows and multiple pairs of trigger  $T$  and de-trigger  $DT$  thresholds were tested (details are in the *Methods* section). Across all tested settings, the detector exhibited limited and unstable performance. With the lowest threshold pair (trigger  $T = 3$ , de-trigger  $DT = 2$ ), the method produced numerous false positives, resulting in a modest accuracy ( $ACC = 0.476 \pm 0.023$ ) and low precision ( $PREC = 0.294 \pm 0.204$ ). Increasing the thresholds to  $T = 4$  and  $DT = 3$  reduced false alarms but also sharply decreased sensitivity ( $SENS = 0.046 \pm 0.095$ ), yielding only marginal improvements in overall accuracy ( $ACC = 0.493 \pm 0.013$ ). Under the most conservative configuration ( $T = 5$ ,  $DT = 3$ ), the detector correctly rejected most non-seismic intervals ( $SPEC = 0.989 \pm 0.025$ ), but its ability to identify true events collapsed ( $SENS = 0.007 \pm 0.018$ ). Because the STA/LTA output is strictly binary and does not generate a continuous detection score, ROC curves and AUC values could not be computed. Overall, the STA/LTA detector performed worse than random guessing and failed to consistently identify seismic-induced perturbations in SOP data. These outcomes (Table 1) indicate that classical amplitude-based detectors are poorly suited to this

sensing modality, reinforcing the need for more robust data-driven approaches.

### Benchmark methods: Fixed-threshold spectral energy analysis

Moreover, we evaluated a second baseline method based on fixed thresholds applied to the band-limited spectral energy of SOP signals. Several combinations of time–frequency parameters were tested, including multiple short time Fourier transform (STFT) window lengths, hop sizes, and alternative band partitions (see *Methods* for further details).

For each configuration, the maximum standardized spectral energy across all bands and channels was used as a univariate detection statistic. Despite this extensive exploration, the detector showed limited ability to identify seismic-induced perturbations. Across all tested settings, the method yielded only moderate accuracy ( $ACC = 0.567 \pm 0.026$ ), with a relatively high precision ( $PREC = 0.714 \pm 0.116$ ) but low sensitivity ( $SENS = 0.264 \pm 0.063$ ). Accordingly, the F1-score remained modest ( $F1 = 0.374 \pm 0.055$ ), indicating that only a minority of seismic events produced spectral-energy excursions sufficiently strong to exceed the fixed threshold. As for the STA/LTA detector, the method produces binary outputs rather than continuous probability scores; therefore, ROC curves and AUC values could not be computed.

Overall, these findings (Table 1) show that fixed spectral-energy thresholds, even when aggregating information across multiple frequency bands and channels, provide only sporadic and inconsistent detections. Notably, this limitation persists even in the most balanced configuration ( $T = 3$ ,  $DT = 2$ ), where the absolute metric values remain low and both sensitivity and specificity exhibit large uncertainties, indicating limited statistical robustness.

### Machine Learning analysis

In this section, we report the results obtained from a binary classification task, aimed at distinguishing seismic events with magnitude  $M \geq 5$  from non-seismic ones. The analysis was carried out using a complete pipeline that included feature engineering and feature selection, followed by a stratified cross-validation procedure with  $k = 10$  folds and  $n = 100$  repetitions. The performance of the models was evaluated using standard classification metrics, including accuracy, precision, recall, F1 score, and area under the ROC curve (AUC). The mean values and standard deviations of each metric, calculated in all cross-validation runs, are reported in Table 2 (Top). The scores obtained by the two ML models employed in this work were statistically different. We found that the distribution of differences in the performance metrics from the two models was approximately Gaussian. By applying a t-test, we verified that the mean of the distribution was significantly different from zero, implying that the outputs of the two models were statistically different. As shown in Table 2 (Top), the XGBoost classifier performed better than Logistic Regression, due to its capability to correctly identify both positive (seismic) and negative (non-seismic) instances. In particular, the XGBoost sensitivity and specificity reach  $\approx 60\%$ , ( $SENS = 0.600 \pm 0.063$  and  $SPEC = 0.592 \pm 0.066$ ). Moreover, the XGBoost classifier returns  $AUC = 0.65 \pm 0.05$ , thus outperforming random guessing. XGBoost confusion matrices and ROC curves are shown in the Supplementary Figure S1. Technical details related to XGBoost and Logistic Regression classifiers are presented in the Supplementary Information (*Supplementary Methods - Machine Learning algorithms* subsection).

To investigate the contribution of individual features to the XGBoost classifier, we conducted a feature-importance analysis using two complementary approaches. First, SHapley Additive exPlanations (SHAP)<sup>21</sup> was implemented to obtain local and global explanations by quantifying the marginal impact of each feature on the model output. Then, the XGBoost's internal tree-based importance metrics were employed to assess the role of variables during model construction. Together, these independent methods provide a robust and cross-validated evaluation of feature relevance.

Figure 2 (a) shows the XGBoost SHAP summary plot, representing the average impact of each feature on the predicted probability of the seismic class, while Figure 2 (b) reports the ranking of the most influential features according to the internal feature–importance measure implemented in XGBoost. For both plots, the results are global, as they are obtained at the end of the stratified and repeated cross-validation procedure ( $k = 10$ ,  $n = 100$ ).

Overall, the two approaches return sets of relevant features that exhibit a pronounced overlap, with only minor rank exchanges between the two representations. The most influential variables are predominantly statistical descriptors associated with robust measures of signal variability, with interquartile ranges playing a major role. Although less frequent, physically motivated features, such as maximum amplitude, spectral roll-off, and entropy, also appear among the relevant predictors. Most of the selected features are derived from the Jones matrix entries; however, among the top-ranked variables, features related to Stokes parameters (e.g., *kurt\_angles\_12*, *kurt\_drift\_47*) are also present. The partial overlap of SHAP values across several features, especially for those ranked lower, suggests that the model decisions rely on combinations of weakly discriminative descriptors rather than on a small set of clearly dominant variables. In this context, the SHAP-based analysis can be interpreted as a form of ablation study, as it estimates the marginal contribution of each feature by comparing model outputs *with and without* the presence of that feature, without requiring explicit retraining-based ablation experiments. This perspective complements the tree-based importance measures and provides an interpretable framework for assessing feature relevance.

The internal XGBoost feature–importance ranking orders features according to their average gain (details in the Supplementary Information, see *Supplementary Methods - Machine Learning algorithms - XGBoost Classifier* subsection), with boxplots

summarizing the distribution of importance values computed across all cross-validation splits ( $counts = n \times k = 1000$ ). The gradual decrease in importance values along the ranking indicates the absence of sharply dominant features, further supporting the idea that the classifier exploits a heterogeneous set of descriptors with comparable relevance.

Finally, in order to quantitatively compare the performance of the baseline models against the best ML classifier (XGBoost), we applied a two-sample Kolmogorov–Smirnov test<sup>22</sup> to sensitivity performances. These tests confirmed that both the benchmark detectors perform worse than XGBoost ( $p < 0.01$ ) for seismic events identification.

### Conditional detectability of seismic events

To quantify the event-wise performance of the ML model and identify potential physical or geometrical factors affecting detectability, we analysed the XGBoost median detection rate (computed over 100 cross-validation repetitions) for each of the 60 earthquakes with  $M \geq 5$ . Figure 3 summarizes the conditional detectability patterns as a function of back-azimuth, epicentral distance, and depth.

Figure 3 (a) shows the median detection rate projected in the back-azimuth–distance domain. Each point corresponds to one earthquake and is coloured according to its detection rate. Most events cluster in the  $180^\circ$ – $225^\circ$  azimuthal range, reflecting the geometry of the regional seismicity relative to the Med-Nautilus cable (endpoints in Catania and Tel Aviv). Within this range, however, the median detection rate exhibits substantial variability: events at similar back-azimuth may display low, intermediate, or high detectability depending on the earthquake. This behaviour indicates that back-azimuth alone does not appear to control detectability, considering that, even at fixed back-azimuth, the detection rate varies with the epicentral distance.

Depth dependence is explored in Figure 3 (b), where events are grouped into four classes determined by the quartiles of the depth distribution. The median detection rate across classes remains relatively stable. Pearson correlation was also computed between depth and detection rate ( $coeff = -0.04$ ,  $p > 0.01$ ), suggesting that, within the depth range spanned by the catalogued events, there is no clear evidence of a systematic dependence of detectability on depth.

Figure 3 (c) shows a detection matrix binned by earthquake magnitude and epicentral distance (10 quantile intervals for distance). Each matrix cell reports the number of events corresponding to specific bins of magnitude and epicentral distance, and is colored according to their median detection rate. Although one might expect detectability to increase with magnitude or decrease with distance, the matrix does not reveal a consistent pattern. Several bins contain only one or two events, and higher-magnitude earthquakes are too few to support generalisable conclusions. As a consequence, no robust relationship between detectability, magnitude, and distance can be inferred from the available data.

### Deep Learning analysis

In this section, we present the results obtained from the DL analysis, which involved the development of an anomaly detection model. We trained an Autoencoder model to reconstruct time series corresponding to the SOP signal, based on two of the three components of the Stokes parameters (the third was excluded since it is determined by the other two, being the sum of the squared components normalized to 1). The training phase was carried out on time series that were theoretically free of seismic-induced anomalies (see *Methods, Machine Learning methods, Target events definition* for further details). Numerical results about the training phase are reported in Table 2 (bottom). After defining an anomaly detection threshold in the training phase (see *Methods, Deep Learning methods, Pre-processing*), the trained Autoencoder was tested on SOP signals that were perturbed by seismic events with magnitude  $M \geq 5$ . In this framework, a large error in the reconstruction of the SOP signal by the autoencoder indicates the occurrence of an anomaly, namely an earthquake. Figure 4 shows the errors of the autoencoder in reconstructing components  $S_1$  and  $S_2$  of the SOP signal for three seismic events:

- Greece, Dodecanese Islands - 31<sup>st</sup>, August, 2022 -  $M = 5.5$ ,  $depth = 10\text{ km}$ ,  $LAT = 37.4783$ ,  $LON = 26.8396$ ,  $distance = 286.8\text{ km}$  (panel (a));
- Greece, Occidental Peloponnese - 29<sup>th</sup>, March, 2024 -  $M = 5.8$ ,  $depth = 33\text{ km}$ ,  $LAT = 37.3764$ ,  $LON = 21.2212$ ,  $distance = 122.4\text{ km}$  (panel (b));
- Turkey - 6<sup>th</sup>, February, 2023 -  $M = 7.9$ ,  $depth = 20\text{ km}$ ,  $LAT = 37.2023$ ,  $LON = 37.0635$ ,  $distance = 604.7\text{ km}$  (panel (c)).

The x-axis in the plots of Figure 4 represents the time window index, with each value corresponding to one of the segments of 10s duration, in which the time series have been divided, with consecutive windows overlapped by 8s. Although the y-axis does not provide a direct time measurement, it condenses the reconstruction errors of the two channels,  $S_1$  and  $S_2$ , into a single mean squared error (MSE) value. Panel (a) of Figure 4 shows a case in which the anomaly detection system does not detect any notable anomalies in the SOP signals ( $S_1$  and  $S_2$ ), despite the actual occurrence of a  $M = 5.5$  seismic event (represented by the red, dotted line in the middle of the plot): the reconstruction error oscillates around the median reconstruction error obtained during Autoencoder training. On the other hand, plots shown in panels (b) and (c) correspond to a successful identification

of SOP signal anomalies, interpretable as perturbations induced by seismic events; in these cases, the reconstruction error suddenly increases after the S-waves arrival. This behavior is particularly evident in the last plot, where the reconstruction error becomes very high. We also attempted to test the system on additional time series of SOP signals that were theoretically expected to exhibit perturbations due to other seismic events with magnitude  $M \geq 5$ . We found that three other earthquakes were reported (see Supplementary Figure S2). However, we observed no other positive outcome: given our anomaly threshold, our system did not detect any further signal anomaly. Technical details related to Autoencoders are presented in the Supplementary Information, in the *Supplementary Methods - Deep Learning algorithms - Autoencoder* subsection.

## Discussion

This study investigates the feasibility of using the SOP of submarine fiber optic cables for the detection of seismic events in the Mediterranean Sea. Although the SOP signals were not originally intended for this purpose, and despite some inherent limitations, we found that seismic events can indeed produce detectable perturbations in the SOP signal under certain conditions. Our results suggest a new avenue for passive seismic monitoring using existing underwater optical infrastructure. After assessing the limitation of benchmark methods (such as STA/LTA and fixed-thresholds spectral energy analyses) in effectively classifying the events in our dataset, we employed two supervised ML models, Logistic Regression and XGBoost, to classify earthquake events based on features extracted from SOP time series. The XGBoost model achieved classification scores around 60% (accuracy, sensitivity, specificity, precision and F1 score), which can appear modest compared to standard ML benchmarks, but are nonetheless remarkable given the noisy and unconventional nature of the signal. Our ML results suggest that certain seismic events can induce patterns in SOP that are distinguishable from background noise; however, the SHAP-based analysis indicates that this distinguishability arises from the cumulative contribution of multiple (largely statistical) features, each providing incremental information, rather than from a small set of clearly dominant predictors.

Moreover, the detectability analysis does not exhibit clear physical or geometric trends exploitable by threshold-based methods. The strong event-to-event variability across all considered parameters indicates that seismic-induced SOP perturbations are not easily separable considering these basic physical quantities, supporting the use of data-driven approaches, capable of combining multiple weak and nonlinear cues beyond classical detection schemes.

To our knowledge, this is the first experimental evidence that real-world SOP signals, as acquired from deployed submarine cables, can encode seismic signatures detectable by ML algorithms.

A previous work by Hernández et al.<sup>23</sup> applied DL methods to distributed acoustic sensing data for earthquake classification, but did not incorporate SOP as a primary feature. More recent studies by Awad et al.<sup>24–26</sup> proposed the use of SOP for early detection in optical networks. However, their methodology relies on simulated SOP signals generated via a tunable waveplate model informed by a limited dataset of real seismic or environmental perturbations. While promising, these works do not provide an experimental validation on raw SOP traces obtained from submarine cables in real operating conditions, as we did in this research. Therefore, our findings constitute a foundational step in the emerging field of SOP-based geophysical monitoring, opening a new research direction that leverages existing optical infrastructure for environmental sensing.

In parallel, we explored an unsupervised DL approach based on an Autoencoder architecture. A comparable method<sup>27</sup> employed a GANN-based architecture leveraging spectrograms derived from the SOP signal. In contrast, our approach operates directly on time series, bypassing the intermediate step of spectrogram extraction. This design choice aims to preserve the temporal structure of the signal and reduce preprocessing overhead, while still enabling effective feature learning.

The goal was to train the model to learn a basic representation of “normal” SOP behavior and identify significant deviations from it as anomalies. This approach has the advantage of requiring minimal pre-processing and no labeled training data, which makes it suitable for near real-time use. However, it is computationally “lazy” in the sense that it does not use domain-specific signal features, nor does it attempt to classify events. Instead, it identifies reconstruction errors above a predefined threshold as potential indicators of seismic activity. Among the 60 recorded earthquakes with  $M \geq 5$  in our dataset, only few events (in particular, a  $M = 5.8$  and a  $M = 7.9$  seismic events) triggered anomalies with significant reconstruction errors. While these few cases indicate that the DL model can recognize meaningful disturbances, the hit rate remains extremely low. Given the limited number of positive detections, it is not currently possible to generalize our results to a broader set of seismic events. We emphasize that the DL anomaly detector is not designed to directly identify earthquakes. Rather, it is intended to highlight significant anomalies in the SOP signal that may be associated with seismic events. Earthquakes with similar geophysical characteristics, such as magnitude, back-azimuth, depth and epicentral distance from the cable, may or may not induce a perturbation in the SOP signal, depending on several factors that are not directly ascribed to these physical quantities. This observation highlights the importance of understanding the coupling efficiency between the cable and the seafloor, as well as the potential heterogeneity of this coupling along the cable’s length. Moreover, the signal-to-noise ratio (SNR) of SOP perturbations is likely influenced by environmental conditions and cable construction details (e.g., cable burial, shielding layers, optical fiber type). Without a detailed physical model of SOP dynamics under mechanical stress, it is difficult to disambiguate true seismic signals from background noise or artifacts induced by the cables. Although generalization remains challenging, the

fact that SOP anomalies associated with major seismic events are detectable in real-world data is a relevant finding itself. Our results offer a convincing proof of concept. By demonstrating that existing optical telecommunication links can be repurposed for seismic detection using SOP analysis, we lay the foundation for a potentially transformative approach to low-cost, large-scale, passive seismic detection. Several directions for future research are immediately apparent. First, further work (feature engineering) is needed to identify and extract richer statistical and spectral descriptors from SOP time series that could improve ML performance. A second promising avenue consists of multi-modal fusion, where SOP data could be integrated with complementary sources such as Distributed Acoustic Sensing (DAS), temperature and strain measurements or external seismic catalogues to improve robustness and detection accuracy. Another important area is the application of transfer learning techniques, particularly those based on domain adaptation, which would allow simulated SOP data, such as that generated by waveplate models, to inform training on real-world measurements. Advancing physical modeling is also crucial: developing a comprehensive mechanical-optical framework for understanding how seismic waves interact with deployed optical fibers would shed light on signal coupling mechanism. Finally, an improved understanding of SOP signal behavior requires systematic noise characterization. We reiterate that the SOP signal, as currently available from telecom-grade submarine fiber links, was never intended as a sensing channel. Its use for seismic detection remains speculative and performance under current conditions is poor. This work represents an initial effort to investigate the use of SOP from submarine fiber-optic cables as a proxy for seismic sensing. Our results suggest that the SOP may contain usable information for detecting strong seismic events ( $M \geq 5$ ). Through a combination of supervised and unsupervised learning methods, we have shown that limited but non-trivial information can be extracted.

Our study serves as a foundational milestone for future research and practical implementations of SOP-based geophysical sensing. Given the global availability of submarine cables, even marginal capabilities could offer a cost-effective supplement to existing seismic infrastructure, especially in oceanic regions where traditional instrumentation is scarce. The outcome of our analyses contributes to bridging this gap by demonstrating that ML/DL models trained on SOP data can achieve accurate and robust detection of offshore seismic events. We validated our results against regional and global earthquake catalogs. This research work highlights the synergy between telecommunication infrastructure, optical sensing and artificial intelligence, revealing a powerful new paradigm for global-scale, cost-efficient seismic monitoring in marine environments. As this field advances, the fusion of photonic sensing with artificial intelligence promises to revolutionize the way we observe and respond to geophysical phenomena across the ocean floor.

## Methods

The objective of this study was the development of an intelligent algorithm for the detection of seismic events based on the analysis of electromagnetic signal variations within a submarine fiber optic cable. In particular, this study focuses on the Med-Nautilus submarine communication cable, which connects countries in the Eastern Mediterranean region. The system comprises 5,729 km of cable and has a total design capacity of 3.84 Tb/s distributed across six fiber optic pairs, representing a strategic asset for the Italian company Telecom Italia Sparkle S.p.A.. The present work focused on the Med-Nautilus section, approximately 2,500 km long, whose endpoints are located in Catania (ITA) and Tel Aviv (ISR), see Figure 5. The electromagnetic signal variations in the fiber optic cable were analyzed in terms of the signal's State of Polarization (SOP). SOP data were provided by Telecom Italia Sparkle S.p.A.: data consist of time series corresponding to the three Stokes parameters and the four entries of the complex Jones matrices, with each entry separated into its real and imaginary components. We analyzed the time series of these quantities over the period from June 1<sup>st</sup>, 2022 to November 30<sup>th</sup>, 2024. To assess the impact of seismic events on the polarization state of the signal in the fiber optic cable, it was necessary to consider also the earthquakes that occurred in the Mediterranean region during the time interval of interest.

Seismic events data were obtained from the website of the Italian National Institute of Geophysics and Volcanology (Istituto Nazionale di Geofisica e Vulcanologia, INGV)<sup>28</sup>. These data were essential for associating potential variations in the SOP signal with one or more seismic events. Following a preprocessing phase of both the SOP signal and the seismic events data, and the application of benchmark detectors (STA/LTA, fixed-threshold spectral energy analyses), we employed ML and DL techniques with a dual objective: (1) to develop an automated binary classification system operating in an offline setting and (2) to implement an anomaly detection system potentially capable of functioning in near real-time.

## Earthquakes dataset creation

Our analysis began with the collection of seismic event data in the Mediterranean region from the INGV portal. These data enabled a precise determination of the number of events together with specific information concerning magnitude, epicentral coordinates and earthquake starting time. For the purposes of our study, we focused on all seismic events with a magnitude  $M \geq 5$ . Subsequently, we estimated the length of the geodesic segment connecting the epicenter of each earthquake to the relevant section of the Med-Nautilus cable. This distance was essential for quantifying the arrival times of primary (P) and secondary (S) seismic waves at the cable. Once these estimates were obtained, it was possible to combine the seismic

information with the time series related to the SOP signal. This association is crucial for the subsequent analyses: seismic events classification by means of both baseline detectors and ML algorithms, and anomaly detection, through a DL Autoencoder.

### Earthquakes data

Seismic event data for the Mediterranean region were downloaded from the INGV website<sup>29</sup>, using the available option “Mediterranean area”. This allows the selection of a rectangular bounding box defined by latitude between  $LAT_{min} = 27$  and  $LAT_{max} = 48$ , and longitude between  $LON_{min} = -7$  and  $LON_{max} = 37.5$ . The period considered for the analysis spans from June 1<sup>st</sup>, 2022 to November 30<sup>th</sup>, 2024. The INGV dataset contains a total of 5851 seismic events with magnitude  $M \geq 2$  on the Richter scale, and provides, for each seismic event, information regarding the latitude and longitude of its epicenter, its starting time, its magnitude and the depth at which it occurred. The number of events decreases as the magnitude  $M$  increases: the exact number of seismic events for different  $M$  intervals is shown in the Supplementary Table S1.

### Epicenter-cable distance and seismic waves velocity estimations

Signal variations in fiber optic cables can be caused by several mechanical perturbations, among which seismic events are of particular relevance. In this study, the detection process was specifically developed for seismic events with magnitude  $M \geq 5$ . The optical fiber signal is sensitive to any kind of mechanical stress, whether directly applied or indirectly induced by thermal fluctuations, along the entire length of the cable. Therefore, in order to be detectable, seismic events must produce signal perturbations that exceed the integrated background noise level over the full cable span. In fact, considering data available to us, the Med-Nautilus cable effectively behaves as a sort of point-like, integral sensor: perturbations occurring at different locations along the cable cannot be spatially distinguished but are instead observed only at the cable endpoint (Catania, ITA), where the SOP signal is recorded. As a consequence, the SOP measurements represent an integrated response of the entire perturbed cable segment rather than a spatially distributed recording.

In this context, the  $M \geq 5$  threshold is well justified and aligns with current scientific literature<sup>16,30–33</sup>. Indeed, the vast majority of seismic events reported as detectable by sensing infrastructures analogous to the Med-Nautilus cable exhibit magnitudes exceeding 5. Documented instances of  $M < 5$  events being detected remain rare; within the datasets investigated in Refs.<sup>16,30–33</sup>, only two seismic events with magnitude below 5 ( $M = 4.4$  and  $M = 4.5$ ) were detected. Notably, these earthquakes were effectively captured because their epicenters were located at a very short distance (on the order of  $\approx 10^1 km$ ) from the Curie submarine fiber cable that revealed them<sup>30</sup>. Consequently, the  $M \geq 5$  threshold represents a balanced and reasonable trade-off for the scope of our analysis.

Seismic events generate distinct wave types<sup>34</sup> that propagate from the hypocenter, the point within the Earth where the earthquake originates, at speeds governed by the depth and the composition of the Earth’s crust and mantle. Consequently, the origin time recorded in the INGV dataset does not coincide with the moment the cable detects the perturbation; this discrepancy arises from the propagation delay over the distance between the seismic source and the fiber. To identify the instant a seismic event alters the SOP signal, we estimated the arrival times of both primary (P) and secondary (S) waves at the Med-Nautilus cable. This required calculating the distances between the epicenter of each earthquake and the cable section of interest. Let  $E$  denote the epicenter of a given earthquake in our dataset. To compute the distance  $d$  between  $E$  and the section of the Med-Nautilus cable under investigation, we adopted the following geometric model:

- We approximated the cable section as the geodesic line passing through Catania (point  $C$ ) and Tel Aviv (point  $T$ ), denoted as  $CT$ . This approximation was justified by the cable’s low curvature and its nearly linear path across the Eastern Mediterranean (see Figure 5).
- We defined our region of interest by tracing two geodesic lines passing through  $C$  and  $T$ , both orthogonal to  $CT$ . The portion of the Mediterranean area lying outside these boundaries was denoted as zone  $S$ .
  - Case 1: if the epicenter  $E$  lay within zone  $S$ , the distance  $d$  was defined as the minimum geodesic distance from  $E$  to either  $C$  or  $T$ .
  - Case 2: if  $E$  lay within the region of interest,  $d$  is computed as the length of the geodesic arc passing through  $E$  and orthogonal to  $CT$ .

It should be noted that this approach is intentionally a simplification of the underlying physics. While the entire cable segment affected by the seismic perturbation undergoes changes in its optical characteristics, the segment closest to the epicenter provides the dominant contribution to the SOP alteration. Furthermore, the complex structure of the Earth’s crust and mantle in the Mediterranean region may cause P- and S- waves to propagate along non-trivial paths. Although further more detailed studies are required to account for these effects, the adopted geometric simplification is considered acceptable and aligns with methodologies used in previous works<sup>31</sup>. As a validation step, we compared a sample of these computed distances

with nominal values provided by an online geodesic distance calculator<sup>35</sup>. The results were consistent, with discrepancies generally on the order of tens of kilometers. Once the distances between the epicenters of the selected earthquakes and the section of the Med-Nautilus cable were calculated, it became possible to estimate the arrival times of P- and S- waves using a parameterized velocity model known as IASP91 (International Association of Seismology and Physics of the Earth's Interior), which is widely adopted in the scientific literature<sup>36</sup>. This model provides estimates of seismic wave velocities as a function of depth. In particular, the IASP91 model is a one-dimensional reference model for the Earth's internal structure, meaning that physical properties vary solely with depth and not laterally. It divides the Earth into several layers: crust, mantle, outer core and inner core. The model accounts for both gradual variations within the Earth and abrupt transitions associated with major discontinuities, such as the Mohorovičić, Gutenberg and Lehmann discontinuities, which respectively separate the outer from the inner mantle, the inner mantle from the outer core and the outer core from the inner core. In addition to the IASP91 model, we also employed a regional, one-dimensional velocity model specific to the Eastern Mediterranean area<sup>37</sup>. By repeating the same procedure described above, we estimated the P- and S-wave arrival times at the Med-Nautilus *CT* cable section for this second model. The resulting travel times were found to be statistically compatible with those obtained using IASP91, showing no significant systematic deviations across the set of considered earthquakes. Technical details and parameterization of the regional model are provided in the Supplementary Information, *Supplementary Methods - Regional 1D velocity model for travel-time estimation* subsection.

Uncertainty in the source–cable distance introduces timing uncertainties in the expected P- and S-wave arrivals, but such uncertainty is unavoidable due to the approximate knowledge of the Med-Nautilus cable geometry. To prevent this from biasing the analysis, we avoided precise arrival-time picking and instead analyzed long post-event windows, starting from the INGV origin time, that fully capture the seismic-induced SOP perturbations. Consequently, timing errors of a few seconds, associated with distance discrepancies of tens of kilometers, are negligible relative to the adopted temporal windows, ensuring robustness to both cable-position uncertainty and complex wave propagation effects. This design choice makes possible to preserve sensitivity to earthquake-induced SOP perturbations.

### State of Polarization (SOP)

The State of Polarization (SOP) of light can be described using the Jones vector  $\tilde{\mathbf{E}}(t)$ , a two-component complex vector representing the amplitude and phase of the electric field components along two orthogonal directions (e.g., the  $x$  and  $y$  axes of a Cartesian coordinate system). Let us consider a monochromatic plane electromagnetic wave propagating in the positive  $z$ -direction of a Cartesian coordinate system, with angular frequency  $\omega$ , wave vector  $k = (0, 0, k_n)$  and wavenumber  $k_n = \omega/c$ , where  $c$  is the speed of light. Formally, the Jones vector is defined as:

$$\tilde{\mathbf{E}}(t) = \begin{pmatrix} E_x(t) \\ E_y(t) \end{pmatrix} = \begin{pmatrix} E_{0x}e^{i(kz-\omega t+\phi_x)} \\ E_{0y}e^{i(kz-\omega t+\phi_y)} \end{pmatrix} = \begin{pmatrix} E_{0x}e^{i\phi_x} \\ E_{0y}e^{i\phi_y} \end{pmatrix} e^{i(kz-\omega t)}, \quad (1)$$

where  $E_{0x}$  and  $E_{0y}$  are the amplitudes of  $\tilde{\mathbf{E}}(t)$  along the  $x$  and  $y$  components,  $\phi_x$  and  $\phi_y$  are the starting phases, and  $t$  represents time. The Jones matrix  $\mathbf{J}$  describes how an optical system modifies a Jones vector  $\tilde{\mathbf{E}}(t)$ , that is, how it alters the polarization state of an electromagnetic signal. A Jones matrix  $\mathbf{J}$  is a  $2 \times 2$  complex matrix that acts on an initial state  $\tilde{\mathbf{E}}_{in}$ , mapping it to a final state  $\tilde{\mathbf{E}}_{fin}$ :

$$\tilde{\mathbf{E}}_{fin} = \mathbf{J}\tilde{\mathbf{E}}_{in}. \quad (2)$$

The polarization state of an electromagnetic signal can also be analyzed in terms of the components of a three-dimensional vector  $\mathbf{S}(t) = (S_1, S_2, S_3)$ , known as Stokes parameters, defined in terms of the  $x$  and  $y$  components of the Jones vector  $\tilde{\mathbf{E}}(t)$ :

$$\begin{aligned} S_1 &= E_x^2 - E_y^2 \\ S_2 &= 2E_xE_y \cos(\delta) \\ S_3 &= 2E_xE_y \sin(\delta) \end{aligned} \quad (3)$$

where  $\delta$  is the angle measured between the phasor  $\tilde{\mathbf{E}}(t)$  and the positive  $x$  axis. In this formalism,  $\tilde{\mathbf{S}}(t)$  is normalized such that

$$|\tilde{\mathbf{S}}(t)| = \sqrt{S_1(t)^2 + S_2(t)^2 + S_3(t)^2} = 1. \quad (4)$$

Given the normalization, it is convenient to represent the Stokes parameters on a three-dimensional unit-radius sphere, known as the Poincaré sphere. A point on this sphere is identified by  $\tilde{\mathbf{S}}(t)$ , which represents the SOP of the electromagnetic signal at a given time instant  $t$ .

## SOP data preprocessing

The analysis of the electromagnetic signal within the Med-Nautilus cable was conducted using an extensive dataset provided by Telecom Italia Sparkle S.p.A. This collection comprises several subsets, consisting of 11 time series that represent the SOP signals recorded at the Catania (ITA) terminal for the Med-Nautilus cable, connecting the endpoints of Catania and Tel Aviv (ISR). Specifically, three of these time series correspond to the components of the Stokes vector  $\mathbf{S}(t)$ , while the remaining eight represent the real and imaginary parts of the complex entries of the Jones matrix  $\mathbf{J}$ , which captures changes in the optical fiber's transmission medium. Each time series covers a duration of approximately three days. Although the signals considered in this study had a nominal sampling frequency of  $f_s \sim 20\text{Hz}$  ( $\Delta t \sim 50\text{ms}$ ), the actual sampling rate exhibited fluctuations on the order of  $1\text{Hz}$ . Furthermore, the data contained discontinuities, likely attributable to hardware malfunctions, maintenance activities or other types of interruptions in the Med-Nautilus system.

To ensure consistency across datasets, the time series were segmented on a daily basis and data points were aggregated every  $200\text{ms}$  using a moving average filter. This procedure yielded uniform daily records with a stabilized sampling frequency of  $f_s = 5\text{Hz}$ ; any remaining data gaps were filled via linear interpolation. Given that the Stokes vector  $\mathbf{S}(t)$  is constrained to the surface of the Poincaré sphere (i.e., its components satisfy the unit-norm condition), the effective degrees of freedom were reduced from 3 to 2. Following a common processing pipeline<sup>16,30,33</sup> we (a) partitioned the signal into non-overlapping time windows of duration  $\Delta t_m = 2\text{s}$ , (b) calculated the mean Stokes vector  $\mathbf{S}_m$  within each window, (c) evaluated the angle  $\theta_m$  between  $\mathbf{S}_m$  and the  $z$ -axis of the Poincaré sphere (corresponding to  $S_3$ ), and (d) rotated  $\mathbf{S}(t)$  ( $t \in \Delta t_m$ ) by  $\theta_m$ . This transformation standardizes the reference frame by aligning the local mean  $\mathbf{S}_m$  within each  $\Delta t_m$  window with a common baseline. By isolating dynamic SOP fluctuations from slow drift, this approach enhances the interpretability of temporal patterns and enables consistent comparisons across different time windows. Figure 6 illustrates an example of this transformation.

## Target events definition

In order to build the pipelines for the subsequent analyses based on benchmark, classical detectors, ML classifier and DL anomaly detector, it was first necessary to distinguish SOP time series associated with seismic activity from those not-associated with seismic events. Consequently, the pre-processed datasets were partitioned into two categories: category *A* comprised days featuring at least one seismic event with a magnitude  $M \geq 5$ , while category *B* included days with either no recorded activity or minor events ( $1 \leq M < 3$ ). This classification yielded 60 datasets for category *A* and 280 datasets for category *B*. To ensure a balanced analysis, we utilized all the 60 datasets from category *A* and the same number of datasets randomly sampled from category *B*. In both cases, diverse complexities must be taken into account. For the category *A* datasets, multiple seismic events may occur in rapid succession. This is exemplified by the seismic swarm of February 6<sup>th</sup>, 2023, which caused numerous casualties in Turkey. On the other hand, the fact that the category *B* datasets are very heterogeneous introduces additional complications. One could argue that it would be more appropriate to consider only datasets related to days entirely free of seismic activity for such a case: however, while datasets without earthquakes do exist (and have been included in the analysis), they are relatively rare. In fact, during the time period considered in this study, the average daily number of recorded seismic events with magnitudes in the range  $2 \leq M \leq 7$  was approximately 6.4.

To facilitate an event-centric analysis, we examined focused segments of the time series rather than full 24-hour recordings. This approach was designed to isolate seismic-induced variations in the SOP signal associated with individual earthquake events. Considering that rupture processes can last several tens of seconds and that signals may be influenced by precursory activity or aftershocks, we defined a time window extending  $15\text{min}$  before and after the event's origin time. This  $30\text{-min}$  interval was empirically selected to capture the main seismic-related signal changes while minimizing contamination from other overlapping seismic events and background noise, thereby enhancing the reliability of the analysis.

To investigate the perturbations induced by P- and S- waves on the SOP signal, we established a target event,  $T_E$ , to serve as a reference within each time series. This event-centric windowing strategy allowed for a focused comparison of SOP signal conditions immediately preceding and following  $T_E$ . definition of  $T_E$  depended on the specific dataset category. For category *A*,  $T_E$  was simply identified as the seismic origin time recorded by INGV on that day. Conversely, identifying  $T_E$  for category *B* was definitely more complex. Since this category may still include low-magnitude events ( $1 \leq M < 3$ ), it was crucial to avoid selecting time series segments potentially affected by actual seismic activity. To ensure this, the following approach was adopted:

- (i) if no seismic events were recorded on the selected day, a random time within the 24-hour period was chosen to be a sort of synthetic  $T_E$ ;
- (ii) if seismic events with magnitude  $1 \leq M < 3$  were present, the 24-hour period was divided into intervals, with boundaries defined by the 00 : 00 of the day, the time(s) of the seismic event(s) and the 24 : 00 of that day. The widest resulting interval,  $I_{max}$ , was then selected. If  $I_{max}$  exceeded  $2\text{h}30\text{min}$  in duration, the synthetic target event  $T_E$  was set  $15\text{min}$  before its upper bound. This choice aimed to maximize the temporal distance from any preceding seismic event and minimize potential contamination.

Once the target event  $T_E$  (either real or synthetic) was established, each time series was portioned into 30-*min* segments centered on this reference time. This process yielded a new set of windowed datasets, labeled as “earthquake YES” or “earthquake NO”. Each sample consisted of 11 time series (specifically, the 3 Stokes parameters and the 8 entries of the Jones matrix) extracted from a window spanning 15 *min* before and after the reference event  $T_E$ .

### SOP data filtering

Once the aforementioned datasets were obtained, a filtering process was required in order to better highlight variations in the SOP signal that could be relevant for the detection of seismic events. This is crucial for the classification analysis based on both benchmark and ML methods. The SOP signal filtering procedure was based on the Power Spectral Density (PSD) technique, implemented using the Welch’s method<sup>38</sup>. Let  $x[n]$  be a discrete-time series, where  $n$  denotes the time index of the sequence, sampled at a frequency  $f_s$ . The Welch’s method for PSD estimation consists of four main steps:

- (a) **Segmentation:** the signal  $x[n]$  is divided into  $K$  segments  $x_k[n]$ , which may partially overlap, each of duration  $L$  seconds. This segmentation helps reduce the variance of the PSD estimate.
- (b) **Windowing:** each segment  $x_k[n]$  is multiplied element-wise by a window function  $w_k[n]$  (e.g., Hamming or Hann window) in a process known as windowing. This step mitigates the effects of truncation and minimizes discontinuities at the segment boundaries.
- (c) **Fourier Transform:** a Discrete Fourier Transform (DFT)  $\mathcal{F}\{\cdot\}$  is computed for each windowed segment  $w_k[n]x_k[n]$ , yielding its frequency-domain representation. For a discrete signal  $x[n]$  of length  $L$ , the DFT is defined as:

$$\mathcal{F}\{x[n]\} = X(f) = \sum_{n=0}^{L-1} x[n] e^{-i2\pi f n/L} \quad (5)$$

or, in the indexed form with respect to the frequency bins  $m$ :

$$X[m] = \sum_{n=0}^{L-1} x[n] e^{-i2\pi \frac{mn}{L}}, \quad m = 0, 1, \dots, L-1. \quad (6)$$

- (d) **PSD computation and averaging:** the PSD  $\hat{S}_x(f)$  is computed for each windowed segment and then averaged across all segments:

$$\hat{S}_x(f) = \frac{1}{K} \sum_{k=1}^K \frac{1}{L} |\mathcal{F}\{w_k[n] \cdot x_k[n]\}|^2. \quad (7)$$

For each dataset of each category (A or B), the PSD was computed on each Stokes parameter over a time span of 15 *min* before and after the target event  $T_E$ . The estimation was performed using overlapping time windows  $w_k[n]$  of duration  $T_f = 100$ s with a 90% overlap. This configuration was selected as a trade-off between frequency resolution and temporal localization. In particular, the window length of  $T_f = 100$ s provides adequate frequency resolution to capture meaningful spectral features potentially linked to seismic-induced perturbations, while ensuring that the signal remains locally stationary within each segment. Furthermore, employing a 90% overlap, following Welch’s method, increases the number of segments available for averaging, thereby reducing the variance of the PSD estimate. This approach is particularly advantageous for our 30-*min* observation windows, as it enhances the statistical reliability of the spectral analysis without significantly increasing computational overhead.

Subsequently, for each of the three components  $S_i$  of the SOP signal, we identified the frequencies  $f_{S_i}$  ( $i = 1, 2, 3$ ) at which the PSD exhibited the greatest increase before and after  $T_E$ . A second-order band-pass Butterworth filter concentrated within a narrow frequency range around  $f_{S_i}$  was applied to each component of the SOP signal. In particular, defining  $\Delta f = \frac{1}{T_f}$  Hz, the frequency interval used for the band-pass filter corresponding to each Stokes parameter  $S_i$  was given by  $[f_{S_i} - \frac{\Delta f}{2}, f_{S_i} + \frac{\Delta f}{2}]$ . The filtering procedure proved particularly effective in highlighting the SOP variations associated with the well-known  $M = 7.9$  earthquake that occurred in Turkey on February 6<sup>th</sup>, 2023. However, this approach failed to emphasize the changes induced by another seismic event ( $M = 6.3$ ) that took place in the same region and on the same day, thereby potentially leading to a false negative in the classification task. Conversely, when examining segments of the SOP signal with no INGV-reported seismic activity, the filter also amplified other fluctuations, presumably attributable to unknown external sources, which could result in false positives.

### Detectors: Benchmark methods

To establish a robust baseline for evaluating the proposed ML approach, we first implemented a set of benchmark methods commonly adopted in classical signal-processing for event detection. These methods rely on predefined rules or thresholding strategies applied to hand-crafted signal attributes, and do not involve data-driven learning or model optimization from labeled examples.

The inclusion of these benchmark methods serves a twofold purpose. First, they provide a well-established reference against which the performance of advanced ML models can be quantitatively assessed, revealing whether the complexity of supervised learning yields measurable improvements. Second, they help contextualize the classification task by clarifying the extent to which physics-inspired or heuristic criteria can exploit the information content of the extracted features.

In this study, benchmark methods are not intended to maximize performance, but rather to serve as reference detectors operating under minimal assumptions and limited adaptivity. Their comparison with ML-based classifiers enables a clearer interpretation of the gains achieved through data-driven modeling, particularly in terms of robustness, generalization capability, and sensitivity to subtle patterns that may not be captured by fixed thresholds or manually designed decision rules.

### STA/LTA

As a first benchmark approach, we implemented the Short-Term Average over Long-Term Average (STA/LTA) analysis on the SOP time series, considering both Stokes parameters and on the real and imaginary components of the Jones matrix. For dataset in categories A and B, we considered a 10-*min* segment starting from the target time  $T_E$ . The short-term and long-term moving averages were computed using the absolute value of the signal via convolution with normalized rectangular windows. To optimize the detector, we explored a range of window durations: the short-term window  $\Delta T_{STA}$  was varied from 1 s to 11 s in steps of 2 s, while the long-term window  $\Delta T_{LTA}$  ranged from 30 s to 150 s in steps of 30 s. Detailed mathematical formulations of these computations are provided in the Supplementary Information, *Supplementary Methods - Benchmark methods: STA/LTA* subsection).

A detection was triggered whenever the STA/LTA ratio exceeded a fixed activation threshold of  $T = 5$ , and remained active until the ratio dropped below a deactivation threshold of  $DT = 3$ . In order to suppress isolated fluctuations, a minimum duration of 1 s above the trigger threshold was required to confirm an event. Multichannel integration was handled as follows: for the SOP data, detections across the two independent stokes components ( $S_1$  and  $S_2$ ) were merged by identifying overlapping activation intervals. For the Jones matrix, the same STA/LTA procedure was applied independently to its eight real and imaginary components ( $JXX_{Re}$ ,  $JXY_{Re}$ ,  $JYX_{Re}$ ,  $JYY_{Re}$ ,  $JXX_{Im}$ ,  $JXY_{Im}$ ,  $JYX_{Im}$ ,  $JYY_{Im}$ ). In this case, an event was confirmed only if at least two components exhibited overlapping triggered intervals. The final output for each time series was a binary indicator, providing a deterministic, rule-based benchmark to quantify the performance gains achieved by the ML and DL models described in the subsequent sections.

### Fixed Spectral Energy Threshold Analysis

A second benchmark was established through a fixed spectral-energy analysis based on the time-frequency representation of the SOP signals. In line with the previous approach, we analyzed a 10-*min* segment starting from the target time  $T_E$  for each dataset in categories A and B. For this analysis, we focused exclusively on the two independent Stokes components ( $S_1$  and  $S_2$ ), processing each channel independently via the short-time Fourier transform (STFT). A grid search was conducted across multiple configurations of window length, hop size, and frequency-band partitioning. Specifically, four different sets of frequency bands were tested:

- Band set 1: (0.01–0.30), (0.30–1.00), (1.00–2.00) Hz,
- Band set 2: (0.01–0.10), (0.10–0.30), (0.30–0.80), (0.80–1.50), (1.50–2.40) Hz,
- Band set 3: (0.01–0.10), (0.10–0.30), (0.30–1.00), (1.00–2.40) Hz,
- Band set 4: (0.01–0.30), (0.30–1.20), (1.20–2.40) Hz.

The selection of these frequency bands was motivated by both physical considerations and established literature. In particular, some of the tested partitions (Band sets 1 and 3) were selected to span broad low-frequency regions (0.01–1 Hz) where long-period surface waves are expected to induce slow SOP perturbations, together with higher-frequency regions (1–2.4 Hz) that may capture more rapid polarization fluctuations. Band set 2 were designed to provide a finer resolution in the 0.01–0.80 Hz range, which has been shown to contain detectable SOP variations during large offshore earthquakes<sup>15,18</sup>. These studies reported clear polarization perturbations associated with low-frequency Rayleigh and Love waves on transoceanic cables, motivating a denser band subdivision in that range. Several bands were therefore chosen in accordance with these earlier findings, while the remaining partitions were intentionally defined in a more exploratory manner to assess whether alternative spectral decompositions could improve seismic-to-noise discrimination. Given the limited prior knowledge on the spectral fingerprints of seismic SOP perturbations on active telecom fibers, this combination of literature-driven and empirical band

definitions was adopted to avoid an overly restrictive parametrization of the spectral-energy detector.

STFT parameters were varied across four window lengths:  $\Delta T_{\text{win}} \in \{20, 40, 80, 120\}$  s, and, for each value of  $\Delta T_{\text{win}}$ , the hop size  $\Delta T_{\text{step}}$  was selected to span target overlaps between 75% and 95%. This yielded hop sizes:

$$\begin{aligned}\Delta T_{\text{step}} &\in \{1, 2, 2.5, 4, 5\} \text{ s} && \text{for } \Delta T_{\text{win}} = 20 \text{ s,} \\ \Delta T_{\text{step}} &\in \{2, 4, 5, 8, 10\} \text{ s} && \text{for } \Delta T_{\text{win}} = 40 \text{ s,} \\ \Delta T_{\text{step}} &\in \{4, 8, 10, 16, 20\} \text{ s} && \text{for } \Delta T_{\text{win}} = 80 \text{ s,} \\ \Delta T_{\text{step}} &\in \{6, 12, 15, 24, 30\} \text{ s} && \text{for } \Delta T_{\text{win}} = 120 \text{ s.}\end{aligned}$$

For each test configuration, the STFT was computed using a Hann window of length  $N_{\text{seg}} = f_s \Delta T_{\text{win}}$ , with overlap  $N_{\text{overlap}} = N_{\text{seg}} - f_s \Delta T_{\text{step}}$ . For every time window and every band  $B_k$ , the spectral energy was computed as:

$$E_k(t) = \sum_{f \in B_k} |S(f, t)|^2 \Delta f,$$

where  $S(f, t)$  denotes the complex-valued STFT of the SOP signal, evaluated at frequency  $f$  and centered at time  $t$  and  $\Delta f$  denotes the median frequency resolution from the STFT. For each channel of the SOP signal, the resulting band–energy time series were transformed via a Box–Cox transformation,  $E_k^{(\lambda)} = \text{BoxCox}(E_k; \lambda_k)$ , with  $\lambda_k$  estimated from the non–earthquake datasets only.  $\lambda_k$  denotes the Box–Cox power parameter associated with the  $k$ -th frequency band, estimated in order to approximate a symmetric, variance-stabilized distribution of the band energy. The transformed energies were scaled across the non–earthquake group using the median and median absolute deviation (MAD):

$$Z_k(t) = \frac{E_k^{(\lambda)}(t) - \text{median}(E_k^{(\lambda)})}{\text{MAD}(E_k^{(\lambda)})}.$$

The final detection score per window was defined as the maximum scaled energy across all bands and all channels:  $Z_{\text{max}}(t) = \max_k Z_k(t)$ . A detection was triggered whenever  $Z_{\text{max}}(t)$  exceeded a fixed activation threshold of  $T = 4$  and the trigger remained active for a minimum of 1s until  $Z_{\text{max}}(t)$  dropped below the deactivation threshold  $DT = 3$ . Continuous intervals above threshold were consolidated into single detections. Similar to the STA/LTA approach, this method provides a deterministic binary classification (event vs. non-event) for each time series and serves exclusively as a baseline for comparison with the more advanced ML and DL techniques described later in this work.

### Detectors: Machine Learning methods

As previously outlined, the primary objective of this study is to develop a robust pipeline for the detection of seismic events on the basis of SOP fluctuations. To this end, a supervised ML approach was implemented to automatically identify potential seismic activity within the signal. Training a supervised ML model requires two key elements:

- a feature matrix  $D$ , where each row represents an individual sample to be classified (in this case, a 30-min SOP segment associated or not-associated with earthquakes) and each column corresponds to a feature, i.e. a statistical or physical descriptor extracted from the preprocessed time series that characterizes the signal’s behavior;
- a set of ground-truth labels, derived from the INGV catalog, which indicate whether the considered SOP segment contains a seismic event (“earthquake yes”) or represents ambient conditions (“earthquake no”).

Consequently, the detection task is framed as a binary classification problem, aimed at distinguishing seismic-induced perturbations from background noise.

### Feature engineering

Following the filtering stage described in the paragraph entitled *SOP data filtering*, we performed feature extraction on the segments of the 11 time series, derived from category A and B datasets. These features are essential for characterizing the signal properties and constructing the training set for the ML models. The proposed approach, which builds on the methodology previously employed by Monaco et al.<sup>39</sup> involves computing a set of statistical and physical features from the 11 available time series:

- Stokes parameters  $S_i$ ,  $i \in \{1, 2, 3\}$ ;

- Jones matrix entries, consisting of eight components  $JXX_{Re}, JXX_{Im}, JXY_{Re}, JXY_{Im}, JYX_{Re}, JYX_{Im}, JYY_{Re}, JYY_{Im}$ , representing the real (Re) and imaginary (Im) parts.

Each time series was segmented as a  $10\text{min}$  window following the target event  $T_E$ . Each of these  $10\text{min}$  segments was further partitioned into blocks of  $10\text{s}$ . Given the sampling frequency of  $f_S = 5\text{Hz}$ , each block consisted of exactly 50 instances. Consequently, each 11-channel segment was divided into 60 blocks. For each block and each channel, we extracted 10 statistical and 20 physical features, spanning both the time and frequency domains. Statistical features, computed in the time domain for each block and for each time series, were:

- Mean  $\mathbf{m}$ ;
- Median  $\mathbf{m}_e$ ;
- Interquartile range  $\mathbf{iqr}$ ;
- Variance  $\sigma^2$ ;
- Variation coefficient  $\frac{\sigma}{\mathbf{m}}$ ;
- Skewness  $\mathbf{s}$ ;
- Kurtosis  $\mathbf{k}$ ;
- Maximum  $\mathbf{max}$ ;
- Minimum  $\mathbf{min}$ ;
- Maximum amplitude  $\mathbf{amp}$ .

Physical features, computed in the time domain for each block of the Stokes parameters time-series, were:

- $\theta_{\mathbf{m}}$ , the angle between mean SOP  $\tilde{\mathbf{S}}_{\mathbf{m}}$  and the  $z$ -axis of the Poincaré sphere;
- Mean, variance, skewness and kurtosis of the distribution of angles  $\theta$  computed between SOP states  $\tilde{\mathbf{S}}$  and mean SOP  $\tilde{\mathbf{S}}_{\mathbf{m}}$ ;
- Mean, variance, skewness and kurtosis of the SOP drift  $\mathbf{v}_s$  distribution;
- **Surf**, the surface occupied by the SOP trajectory on the Poincaré sphere;
- $\mathbf{A}_{\mathbf{S}i}$ ,  $i \in \{1, 2, 3\}$ , which is the area under the PSD curve computed for each Stokes parameter.

The remaining physical features were computed for each block and for each time-series. They were:

- Energy  $\mathbf{E}$ ;
- Power  $\mathbf{P}$ ;
- Zero crossing rate  $\mathbf{ZCR}$ ;
- Entropy  $\mathbf{H}$ ;
- Dynamic Time Warping distance  $\mathbf{DWT}_d$ ,

defined in the time domain, and

- Spectral centroid  $\mathbf{C}_s$ ;
- Spectral entropy  $\mathbf{H}_s$ ;
- Spectral flux  $\mathbf{F}_s$ ;
- Spectral roll-off  $\mathbf{R}_s$ ,

defined in the frequency domain. Below are the non-trivial formal definitions for the primary physical features in the time domain.

- **SOP drift  $v_s$**  quantifies the rate of change of the Stokes vector, which represents the SOP of the signal on the Poincaré sphere:

$$v_s = \sqrt{\sum_{i=1}^3 \left( \frac{dS_i}{dt} \right)^2}. \quad (8)$$

- **Energy  $E$** . Let  $x_i(n)$ ,  $n = 1, \dots, N$ , the  $n$ -th instance of the  $i$ -th block related to the generic time-series  $x$ . Then, energy  $E$  quantifies the total signal magnitude over a time window:

$$E = \sum_{n=1}^N |x_i(n)|^2. \quad (9)$$

- **Power  $P$**  quantifies the average energy per sample within a signal segment:

$$P = \frac{1}{N} \sum_{n=1}^N |x_i(n)|^2. \quad (10)$$

- **Zero crossing rate  $ZCR$**  measures the frequency of sign changes in the signal:

$$ZCR = \frac{1}{2N} \sum_{i=1}^N |sgn[x_i(n)] - sgn[x_i(n-1)]|, \quad (11)$$

where  $sgn(k) = \pm 1$  when  $k \gtrless 0$ .

- **Entropy  $H$**  quantifies the signal complexity and uncertainty:

$$H = - \sum_{j=1}^K e_j \log_2(e_j), \quad (12)$$

where  $e_j$  is the ratio between the energy of the  $j$ -th frame (among the  $K$  frames into which the  $i$ -th chunk is divided) and the total energy of that  $i$ -th chunk.

- **Dynamic Time Warping (DTW)**. Given two time series  $X = (x_1, x_2, \dots, x_N)$  and  $Y = (y_1, y_2, \dots, y_M)$ ,  $DTW$  computes their similarity by finding the optimal alignment path that minimizes the cumulative distance, allowing for temporal stretching and compression. The local distance between two elements is typically defined using the squared Euclidean distance  $d(x_i, y_j) = (x_i - y_j)^2$ . The cumulative  $DTW$  distance is given by:

$$DTW(X, Y) = \min_W \sum_{k=1}^K d(w_k), \quad (13)$$

where  $W = (w_1, w_2, \dots, w_K)$  is a warping path that aligns the elements of  $X$  and  $Y$  so that  $d(w_k) = d(x_{i_k}, y_{j_k})$ . This path must satisfy the following constraints:

- **Monotonicity**: the path cannot go backwards in time;
- **Continuity**: the path must not skip elements;
- **Boundary conditions**: the path must start at  $(1, 1)$  and end at  $(N, M)$ .

The following recursive algorithm was used to compute the optimal alignment  $D$  between adjacent blocks of the SOP signal:

$$D(i, j) = d(x_i, y_j) + \min \{D(i-1, j), D(i, j-1), D(i-1, j-1)\}. \quad (14)$$

Let  $X_i(n)$ ,  $n = 1, \dots, N$  be the Fourier coefficients obtained applying the DFT on the  $i$ -th chunk. The frequency-domain features employed in this study are defined as follows.

- **Spectral Centroid  $C_s$**  indicates the average position of a frequency distribution, weighted by the intensity (or amplitude) values associated with each frequency:

$$C_s = \frac{\sum_{n=1}^{N/2} n X_i(n)}{\sum_{n=1}^{N/2} X_i(n)}. \quad (15)$$

- **Spectral Entropy  $H_s$**  is the frequency-domain equivalent of entropy  $H$ :

$$H_s = - \sum_{j=1}^K n_j \log_2(n_j) \quad (16)$$

- **Spectral Flux  $F_s$**  measures the change in the normalized spectral energy distribution  $EN_i(n)$  between consecutive windows:

$$F_s(i, i-1) = \sum_{n=1}^{N/2} (EN_i(n) - EN_{i-1}(n))^2, \quad (17)$$

with

$$EN_i(n) = \frac{X_i(n)}{\sum_{n=1}^{N/2} X_i(n)}. \quad (18)$$

- **Spectral Roll-off  $R_s$**  describes the rate at which the spectral intensity or amplitude of a signal decreases as frequency increases:

$$R_s = C \sum_{n=1}^{N/2} X_i(n), \quad (19)$$

where  $C$  is a threshold frequency below which most (typically 90%) of the spectral amplitude distribution is concentrated.

The process of extracting all the statistical and physical features introduced so far was repeated on each 10-*min* segment obtained from the 11 time series previously identified in datasets of categories A and B. The resulting global dataset  $D_f$  has a size of  $120 \times 13,197$ . Of the 120 rows, 60 correspond to time series segments belonging to category A datasets; the remaining 60 rows correspond to instances belonging to category B datasets.  $D_f$  columns represent the collection of statistical and physical features computed on each chunk into which each time series was portioned. This final dataset served as the input for the subsequent seismic classification.

### Feature selection

At the conclusion of the feature extraction process, a feature selection analysis was undertaken. In the context of high-dimensional data analysis, such as in the present case study, feature selection is a critical step to enhance the performance of ML predictive models, reduce noise and improve computational efficiency. The global dataset  $D_f$  consisted of 120 rows, representing samples of seismic and non-seismic events, with a number of features slightly exceeding 13,000. A careful selection of the most relevant features was necessary, aiming to minimize overfitting and thereby improve the generalization capability of the artificial intelligence model. The adopted approach combines two methods: a variance analysis and a Pearson correlation analysis (which evaluates linear correlation) among the starting  $\simeq 13,000$  features. Following this dual analysis, redundant and minimally informative variables were removed. Regarding the variance analysis, the values of each feature  $X$  were temporarily scaled to the range  $[0, 1]$  using the Min-Max Scaling method:

$$X' = \frac{X - X_{\min}}{X_{\max} - X_{\min}}, \quad (20)$$

where  $X_{\min}$  is the minimum value of the feature,  $X_{\max}$  is the maximum value, and  $X'$  is the normalized value scaled to the range  $[0, 1]$ . This normalization of features to the range  $[0, 1]$  prior to the variance analysis is fundamental, since features with smaller numerical scales could exhibit lower variance merely due to the magnitude of their values and not because they are truly less informative. After normalization, all of the over 13,000 features had distributions with values confined within the interval  $[0, 1]$ .

It should be noted that this type of transformation is linear and does not alter the shape of the distribution of each feature. At this point, the variance of each feature was calculated, defined as

$$\sigma^2 = \frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2, \quad (21)$$

where  $N$  is the number of instances,  $x_i$  is the  $i$ -th value of the feature,  $\bar{x}$  is the mean value of the feature. All features whose variance was below the value corresponding to the 85<sup>th</sup> percentile of the distribution ( $\sigma^2 \approx 0.06$ ) were eliminated. The vast majority of features exhibited very low variance, suggesting a lack of relevant informative content within them. Hence, the variance analysis enabled a reduction in the number of features from over 13,000 to less than 2,000.

Following the variance-based feature selection, a correlation analysis was performed with the aim of removing redundant features, i.e., those characterized by the same informative content. The presence of redundant, and thus highly correlated, features can negatively impact the effectiveness of predictive models by increasing the risk of overfitting. The correlation between two features  $X$  and  $Y$  was calculated using the well-known Pearson correlation coefficient, defined by the formula:

$$\rho_{X,Y} = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2 \sum_{i=1}^N (y_i - \bar{y})^2}}. \quad (22)$$

The Pearson correlation coefficient  $\rho_{X,Y}$  ranges in  $[-1, 1]$ . For the purposes of the analysis, it was not necessary to distinguish between positive and negative correlation, but only to identify pairs of features with a strong mutual dependence. Therefore, the absolute value of the correlation coefficient  $|\rho_{X,Y}|$  was considered. An elimination threshold of  $\rho_t = 0.9$  was set, as such a threshold is commonly considered in the literature to indicate a high correlation condition. The Pearson correlation was computed for every pair of features. For each identified pair, one of the two features was randomly removed only if  $|\rho_{X,Y}|$  exceeded the threshold  $\rho_t$ . This approach considerably reduced possible informational redundancy. Random elimination was chosen to avoid imposing arbitrary priorities among features with similar informational content. The correlation analysis further reduced the cardinality of the dataset  $D_f$ , preserving a final subset of 165 features. These features were thus employed within the ML pipeline.

### Training and testing the Machine Learning models

The final dataset consists of 120 instances, equally divided between seismic and non-seismic events (60 instances per class), and 165 features, corresponding to the most informative variables retained after dimensionality reduction via variance analysis and redundancy removal through Pearson correlation. To evaluate the predictive performance of ML models on this dataset, we implemented a pipeline based on stratified  $k$ -fold cross-validation with  $k = 10$ , repeated  $n = 100$  times. Stratified cross-validation ensures that each fold maintains the original class distribution, which is particularly important in binary classification tasks to prevent biased training or evaluation due to imbalanced splits. Within each step of the 10-fold pipeline, the following were conducted:

1. The data were split into a training set, including 9 out of the 10 stratified folds (i.e. 90% of the instances), and a test set, consisting of the remaining one (10% of the dataset);
2. The feature distributions in the training set were standardized using z-score normalization (zero mean and unit variance), and the same transformation parameters were applied to standardize features in the test set;
3. Two ML models were fitted on the training set: Logistic Regression, used as a benchmark, and XGBoost;
4. Predictions were generated for both models on the test set.

The predictions obtained at each fold were stored and, only at the end of an entire cross-validation cycle (i.e., after repeating the pipeline on all the 10 folds), the evaluation metrics were computed. The performance was assessed in terms of accuracy, recall, precision, specificity, and  $F1$ -score, providing a comprehensive view of the classification quality across all repetitions.

### Detectors: Deep Learning methods

As previously stated at the beginning of the *Methods* section, the aim of this part of the analysis is to implement an anomaly detection system capable of operating in near real-time. The purpose of the system is to flag the possible presence of significant anomalies, which may also coincide with seismic events. The core of the system consists of an Autoencoder whose architecture integrates both Multilayer Perceptron (MLP) and Long Short-Term Memory (LSTM) networks, employed in both the encoder and decoder stages. This hybrid structure enables the model to capture nonlinear relationships as well as temporal dependencies

within the input time series. The Autoencoder is trained exclusively on time series belonging to category  $B$ , i.e. data that are as free from anomalies as possible. This allows the model to learn and reconstruct the normal behavior of the system. During the testing phase, many time series belonging to category  $A$  (containing seismic anomalies characterized by  $M \geq 5$ ) are passed through the trained Autoencoder; significant reconstruction errors indicate deviations from the learned patterns, thus signaling the presence of anomalous events.

### Pre-processing

The training of the Autoencoder started considering the 280 datasets belonging to category  $B$ , each containing 11 time series (comprising the Stokes parameters and the entries of the Jones matrix). Let us recall that each of these datasets contains time series of  $30\text{min}$  in length, centered around the target event  $T_E$ ; in the case of category  $B$ ,  $T_E$  is a fictitious event. As a first step, we excluded the signals related to the entries of the Jones matrix in order to implement a simpler and more streamlined processing pipeline. With regard to the Stokes parameters, we focused exclusively on the  $S_1$  and  $S_2$  components, having previously rotated the vectors to constrain the  $S_3$  component to exhibit minimal oscillations around 1. This transformation allowed us to omit  $S_3$  from the analysis without any relevant loss of information. Up to this point, we have assembled 280 datasets from category  $B$ , each containing two channels ( $S_1$  and  $S_2$ ), with time series of  $30\text{min}$  in length centered around the fictitious target event  $T_E$ . These constitute the training dataset for the anomaly detection system we developed.

### Training and testing the Autoencoder

The category  $B$  datasets contain two-channel time series,  $S_1$  and  $S_2$ , each of which is  $\Delta t = 30\text{min}$  long. Let  $S_i$  denote the tensor containing a pair of time series ( $S_1, S_2$ ) from the  $i$ -th dataset, with  $i \in \{1, \dots, 280\}$ . Each of these time series pairs is segmented into shorter sub-sequences of  $\delta t_{sub} = 10\text{s}$  in duration, with an overlap of  $\delta t_{overlap} = 8\text{s}$  and a stride of  $\delta t_{stride} = 2\text{s}$ . The total number of sub-sequences is given by  $N = \left\lfloor \frac{\Delta t - \delta t_{sub}}{\delta t_{stride}} \right\rfloor + 1 = 896$ . We denote each sub-sequence as  $s_{ij,c}$ , with  $j \in \{1, \dots, 896\}$  and  $c \in \{1, 2\}$ . Considering the full set of 280 datasets, the total number of sub-sequences  $s_{ij,c}$  is 250,880. Each  $s_{ij,c}$  is a pair of  $\delta t_{sub} = 10\text{s}$  time series and represents the input type that the Autoencoder is trained to reconstruct. A figurative example of the  $s_{ij,c}$  sample is shown in Figure 7. To train the Autoencoder effectively, it is appropriate to apply z-score standardization. We applied z-score standardization to the tensor obtained by concatenating all the sub-sequences  $s_{ij,c}$ . This normalization involves computing a mean  $\mu_c$  and a standard deviation  $\sigma_c$  for each channel  $c$ : the scaling parameters must learn across the entire training set, so that they can subsequently be applied to the test data without causing data leakage. Out of the total number of input samples, 90% were used to train the Autoencoder, while the remaining 10% were used for validation. The training of the Autoencoder was performed over a maximum of 200 epochs, with early stopping triggered in absence of improvements for 10 consecutive epochs. Training and validation losses and general performance of the Autoencoder are shown in Supplementary Figure S2. The structure of the Autoencoder employed in this analysis is described in the dedicated *Supplementary Methods - Deep Learning algorithms - Autoencoder* subsection of the Supplementary Information. Since the Autoencoder is intended to operate as an anomaly detector, it is necessary to define a threshold for anomaly classification,  $\eta$ . To this end, we analyzed the distribution of the reconstruction error on the training set. Once the Autoencoder model was trained, the sub-sequences  $s_{ij,c}$  were fed back into the network as input and a reconstruction error  $\text{MSE}_{ij}$  was computed for each of them. The reconstruction error  $\text{MSE}_{ij}$  for a single sub-sequence  $s_{ij,c}$  is defined as:

$$\text{MSE}_{ij} = \frac{1}{\delta t_{sub} \cdot C} \sum_{t=1}^{\delta t_{sub}} \sum_{c=1}^C (x_{ij,t,c} - \hat{x}_{ij,t,c})^2 \quad (23)$$

where:

- $x_{ij,t,c}$  denotes the true value of the signal at time  $t$  and for channel  $c$  ( $c \in \{1, 2\}$ ) for the sub-sequence  $s_{ij,c}$ ;
- $\hat{x}_{ij,t,c}$  is the corresponding reconstructed value;
- $C$  is the total number of channels,  $C = 2$ ;
- $\delta t_{sub}$  is the length of the sub-sequence (in number of time steps).

Each dataset  $S_i$  yields 896 values of  $\text{MSE}_{ij}$  and, considering all 280 training datasets, the total number of  $\text{MSE}_{ij}$  values is 250,880. The anomaly threshold  $\eta$  was set based on the distribution of the reconstruction errors in the training set; in particular, it was set equals to the upper inner fence of the  $\text{MSE}_{ij}$  distribution:  $\eta = Q_3 + 3/2 \cdot \text{IQR}$ , where  $Q_3$  is the third quartile of the  $\text{MSE}_{ij}$  distribution and  $\text{IQR} = Q_3 - Q_1$  is the interquartile range. Figure 8 shows the  $\text{MSE}_{ij}$  distribution:  $\eta \approx 0.19$ . It is important to note that the reconstruction error does not evolve as a continuous-time signal. Instead, it is evaluated independently for each sub-sequence  $s_{ij,c}$ , yielding a discrete sequence of reconstruction error values  $\text{MSE}_{ij}$  computed every  $\delta t_{stride} = 2\text{s}$ . Each

$MSE_{ij}$  summarizes the reconstruction quality over a window of duration  $\delta t_{sub} = 10s$ , rather than representing an instantaneous measurement. As a consequence, isolated threshold crossings of the reconstruction error may arise from short-lived fluctuations and do not necessarily correspond to physically meaningful anomalies. To identify persistent anomalies in this discrete-time error sequence, we implemented a trigger–detrigger ( $T - DT$ ) mechanism initialized at the P-wave arrival. The system is activated when the reconstruction error exceeds the trigger threshold  $T = \eta$  and remains active as long as the error stays above the detrigger threshold  $DT = Q_3$ . Only anomaly segments persisting for at least five consecutive sub-sequences are retained. Given the window length and overlap, this criterion corresponds to detecting events with a minimum effective duration of approximately 10s.

In the testing phase, we evaluated the trained Autoencoder on a restricted subset of datasets from category A. Finally, the two time series of each dataset  $S_k$  were standardized using the parameters  $\mu_c$  and  $\sigma_c$ ,  $c \in \{1, 2\}$ , computed during the training phase. Each standardized time series was then segmented into sub-sequences  $s_{km,c}$ , following the same procedure used during training. These sub-sequences were passed through the trained Autoencoder, which returned a reconstruction error  $MSE_{km}$  for each of them. The reconstruction error obtained for some examined cases is shown in Figure 4.

## Data Availability

The datasets generated and/or analysed during the current study are publicly available in a dedicated *Zenodo* repository<sup>40</sup>. The repository includes the processed State of Polarization (SOP) time series, metadata describing seismic events, and the derived feature matrix used for Machine Learning analyses.

## Code Availability

The code used for data preprocessing, feature extraction, model training, and evaluation is publicly available in the aforementioned *Zenodo* repository<sup>40</sup>. The repository includes scripts for spectral filtering, feature engineering, supervised and unsupervised machine learning models, and figure generation.

## References

1. Cremen G., Galasso C., Zuccolo E. Investigating the potential effectiveness of earthquake early warning across Europe. *Nat Commun* **13** (2022).
2. Ladina C., Marzorati S., Amato A., Cattaneo M. Feasibility Study of an Earthquake Early Warning System in Eastern Central Italy. *Front. Earth Sci.* **9** (2021).
3. Cervone G., Kafatos M., Napolitano D., Singh R. P. An early warning system for coastal earthquakes. *Adv. Space Res.* **37**, 636–642 (2006).
4. Allen R. M., Gasparini P., Kamigaichi O., Bose M. The Status of Earthquake Early Warning around the World: An Introductory Overview. *Seismol. Res. Lett.* **80(5)**, 682–693 (2009).
5. Allen R. M., Melgar D. Earthquake Early Warning: Advances, Scientific Challenges, and Societal Needs. *Annu. Rev. Earth Planet. Sci.* **47**, 361–388 (2019).
6. Kuyuk H. S., Allen R. M. Optimal Seismic Network Density for Earthquake Early Warning: A Case Study from California. *Seismol. Res. Lett.* **84(6)**, 946–954 (2013).
7. Lindsey N. J., Martin E. R. Fiber-Optic Seismology. *Annu. Rev. Earth Planet. Sci.* **49**, 309–336 (2021).
8. Mecozzi A. Sensing with submarine optical cables. *APL Photonics* **9(7)** (2024).
9. Yu J., Xu P., Yu Z., Wen K., Yang J., Wang Y., Qin Y. Principles and Applications of Seismic Monitoring Based on Submarine Optical Cable. *Sensors* **23** (2023).
10. Matias L., Carrilho F., Sa V., Omira R., Nieheus M., Corela C., Barros J., Omar Y. The Contribution of Submarine Optical Fiber Telecom Cables to the Monitoring of Earthquakes and Tsunamis in the NE Atlantic. *Front. Earth Sci.* **9** (2021).
11. Marra G. et al. Optical interferometry–based array of seafloor environmental sensors using a transoceanic submarine cable. *Sci.* 874–879 (2022).
12. Williams E. F., Fernández-Ruiz M.R., Magalhaes R. et al. Distributed sensing of microseisms and teleseisms with submarine dark fibers. *Nat Commun* **10** (2019).
13. Idrissi A., Bonanno D., Di Mauro L. S., Diego-Tortosa D., Gomez-Garcia C., Ker S., Le Pape F., Murphy S., Pulvirenti S., Riccobene G. et al. Detection of Seismic and Acoustic Sources Using Distributed Acoustic Sensing Technology in the Gulf of Catania. *J. Mar. Sci. Eng.* **13** (2025).

14. Marra G., Clivati C., Xuereb A. et al. Ultrastable laser interferometry for earthquake detection with submarine optical cables. *Sci.* **361(6401)**, 486–490 (2018).
15. Costa L., Varughese S., Mertz P. et al. Localization of seismic waves with submarine fiber optics using polarization-only measurements. *Commun Eng* **2** (2023).
16. Zhan Z., Cantono M., Kamalov V., Mecozzi A., Müller R., Yin S., Castellanos J. C. Optical polarization-based seismic and water wave sensing on transoceanic cables. *Sci.* **371**, 931–936 (2021).
17. Sladen A., Rivet D., Ampuero J.P. et al. Distributed sensing of earthquakes and ocean-solid Earth interactions on seafloor telecom cables. *Nat. Commun.* **10** (2019).
18. Zhan, Z. Ocean-Bottom Seismology Using Optical Fiber: A Global Opportunity. *Bull. Seismol. Soc. Am.* **110(6)**, 2925–2934 (2020).
19. Perol T., Gharbi M., Denolle M. Convolutional neural network for earthquake detection and location. *Sci. Adv.* **4(2)** (2018).
20. Mousavi S. M., Ellsworth W. L., Zhu W. et al. Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking. *Nat. Commun.* **11(1)** (2020).
21. Lundberg S., Lee S. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* **30** (2017).
22. Smirnov, N. Table for estimating the goodness of fit of empirical distributions. *The Annals Math. Stat.* **19**, 279–281 (1948).
23. Hernández P. D., Ramírez J. A., Soto M. A. Deep-Learning-Based Earthquake Detection for Fiber-Optic Distributed Acoustic Sensing. *J. Light. Technol.* **40**, 2639–2650 (2022).
24. Awad H., Usmani F., Virgillito E., Bratovich R., Proietti R., Straullu S., Pastorelli R., Curri V. Seismic detection through state-of-polarization analysis in optical fiber networks. *Proc. SPIE 12835, Opt. Fibers Sensors for Med. Diagn. Treat. Environ. Appl.* (2024).
25. Awad H., Straullu S., Usmani F., Virgillito E., Bratovich R., Aquilino F., Proietti R., Curri V. Experimental earthquake early detection through polarization changes in intelligent optical networks. *Proc. SPIE 12835, Opt. Fibers Sensors for Med. Diagn. Treat. Environ. Appl.* (2025).
26. Awad H., Usmani F., Virgillito E., Bratovich R., Straullu S., Aquilino F. Artificial Intelligence Driven Earthquakes Early Detection in Noisy Urban Areas. *IEEE Middle East Conf. on Commun. Netw. (MECOM), Abu Dhabi, UAE* 65–70 (2024).
27. Abdelli K., Lonardi M., Gripp J., Olsson S., Boitier F., Layec P. Unsupervised Anomaly Detection and Localization with Generative Adversarial Networks. *ECOC 2024, 50th Eur. Conf. on Opt. Commun. Frankfurt, Ger.* 998–1001 (2024).
28. Istituto Nazionale di Geofisica e Vulcanologia. <https://www.ingv.it/en/>. Accessed: 2025-07-21.
29. Istituto Nazionale Geofisica e Vulcanologia. Earthquakes list. <https://terremoti.ingv.it/>. Accessed: 2025-01-07.
30. Cantono M., Kamalov V., Vusirikala V., Salsi M., Newland M., Zhan Z. Sub-Hertz Spectral Analysis of Polarization of Light in a Transcontinental Submarine Cable. *2020 European Conference on Optical Communications, Brussels, Belgium* (2020).
31. Castellanos J. C., Zhan Z., Kamalov V., Cantono M., Yin S., Mecozzi A., Bhattacharya S., Allen R. M. Optical polarization-based sensing and localization of submarine earthquakes. *Opt. Fiber Commun. Conf. Exhib. San Diego, CA, USA* 1–3 (2022).
32. Cantono M., Castellanos J. C., Bhattacharya S., Yin S., Zhan Z., Mecozzi A., Kamalov V. Optical Network Sensing: Opportunities and Challenges. *Opt. Fiber Commun. Conf. Exhib. San Diego, CA, USA* (2022).
33. Mecozzi A., Cantono M., Castellanos J. C., Kamalov V., Müller R., Zhan Z. Polarization sensing using submarine optical cables. *Opt.* **8**, 788–795 (2021).
34. Honda H. Earthquake Mechanism and Seismic Waves. *J. Phys. Earth* **10**, 1–97 (1962).
35. Distance Calculator. <https://en.distance.to/>. Accessed: 2025-07-21.
36. Kennett, B. L. N., Engdahl, E. R. Traveltimes for global earthquake location and phase identification. *Geophys. J. Int.* **105**, 429–465 (1991).
37. Blom N., Gokhberg A., Fichtner A. Seismic waveform tomography of the central and eastern mediterranean upper mantle. *Solid Earth* **11**, 669–690 (2020).

38. Welch P. D. The use of Fast Fourier Transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio Electroacoustics* 70–73 (1967).
39. Monaco A., Amoroso N., Bellantuono L., Pantaleo E., Tangaro S., Bellotti R. Multi-Time-Scale Features for Accurate Respiratory Sound Classification. *Appl. Sci.* **10** (2020).
40. Caruso M., Morelli M. Seismic detection using submarine cable polarization signals with machine learning. <https://doi.org/10.5281/zenodo.18774527>. Accessed: 2025-02-25.

## Acknowledgements

This study received no external funding.

## Author contributions

Mario Caruso data analysis, data curation, original draft writing; Michele Morelli data analysis, data curation, original draft writing; Alfonso Monaco supervision, original draft writing and draft review; Nicola Amoroso draft review; Luigi Avagliano draft review; Tommaso Maggipinto draft review; Danilo Decaroli draft review; Alberto Marullo draft review; Pietro Patimisco draft review; Angelo Sampaolo draft review; Francesco Simeone data analysis, draft review; Vincenzo Spagnolo draft review; Andrea Zifarelli draft review; Loredana Bellantuono supervision, original draft writing and draft review; Roberto Bellotti project management, supervision and draft review.

## Competing interests

The authors declare no competing interests.

Benchmark Models	ACC	black SENS	SPEC	PREC	F1
STA/LTA (T = 3 / DT = 2)	0.476 ± 0.023	0.216 ± 0.264	0.736 ± 0.301	0.294 ± 0.204	0.209 ± 0.221
STA/LTA (T = 4 / DT = 3)	0.493 ± 0.013	0.046 ± 0.095	0.939 ± 0.106	0.128 ± 0.201	0.062 ± 0.118
STA/LTA (T = 5 / DT = 3)	0.498 ± 0.005	0.007 ± 0.018	0.989 ± 0.025	0.054 ± 0.140	0.012 ± 0.032
Spectral-energy (T = 3 / DT = 2)	0.547 ± 0.024	0.562 ± 0.142	0.491 ± 0.167	0.549 ± 0.031	0.565 ± 0.055
Spectral-energy (T = 4 / DT = 3)	0.567 ± 0.026	0.264 ± 0.063	0.870 ± 0.096	0.714 ± 0.116	0.374 ± 0.055
Spectral-energy (T = 5 / DT = 3)	0.550 ± 0.007	0.115 ± 0.019	0.984 ± 0.019	0.905 ± 0.106	0.203 ± 0.027

**Table 1. Benchmark performances of simple threshold-based detectors applied to State of Polarization signals.**

**Top:** Detection results for the STA/LTA (Short Time Average/Long Time Average) method across three pairs of trigger  $T$  and de-trigger  $DT$  thresholds.

**Bottom:** Performance of the fixed-threshold spectral energy detector, evaluated across multiple window sizes, hop sizes, and frequency bands.

In all cases sensitivity remains low, confirming that simple amplitude- or energy-based detectors struggle to capture seismic-induced State of Polarization perturbations.

ML Models	ACC	SENS	black SPEC	PREC	F1	AUC
XGB	0.596 ± 0.040	0.600 ± 0.063	0.592 ± 0.066	0.596 ± 0.042	0.596 ± 0.045	0.651 ± 0.047
LogR	0.561 ± 0.034	0.578 ± 0.046	0.551 ± 0.060	0.563 ± 0.040	0.570 ± 0.035	0.586 ± 0.038

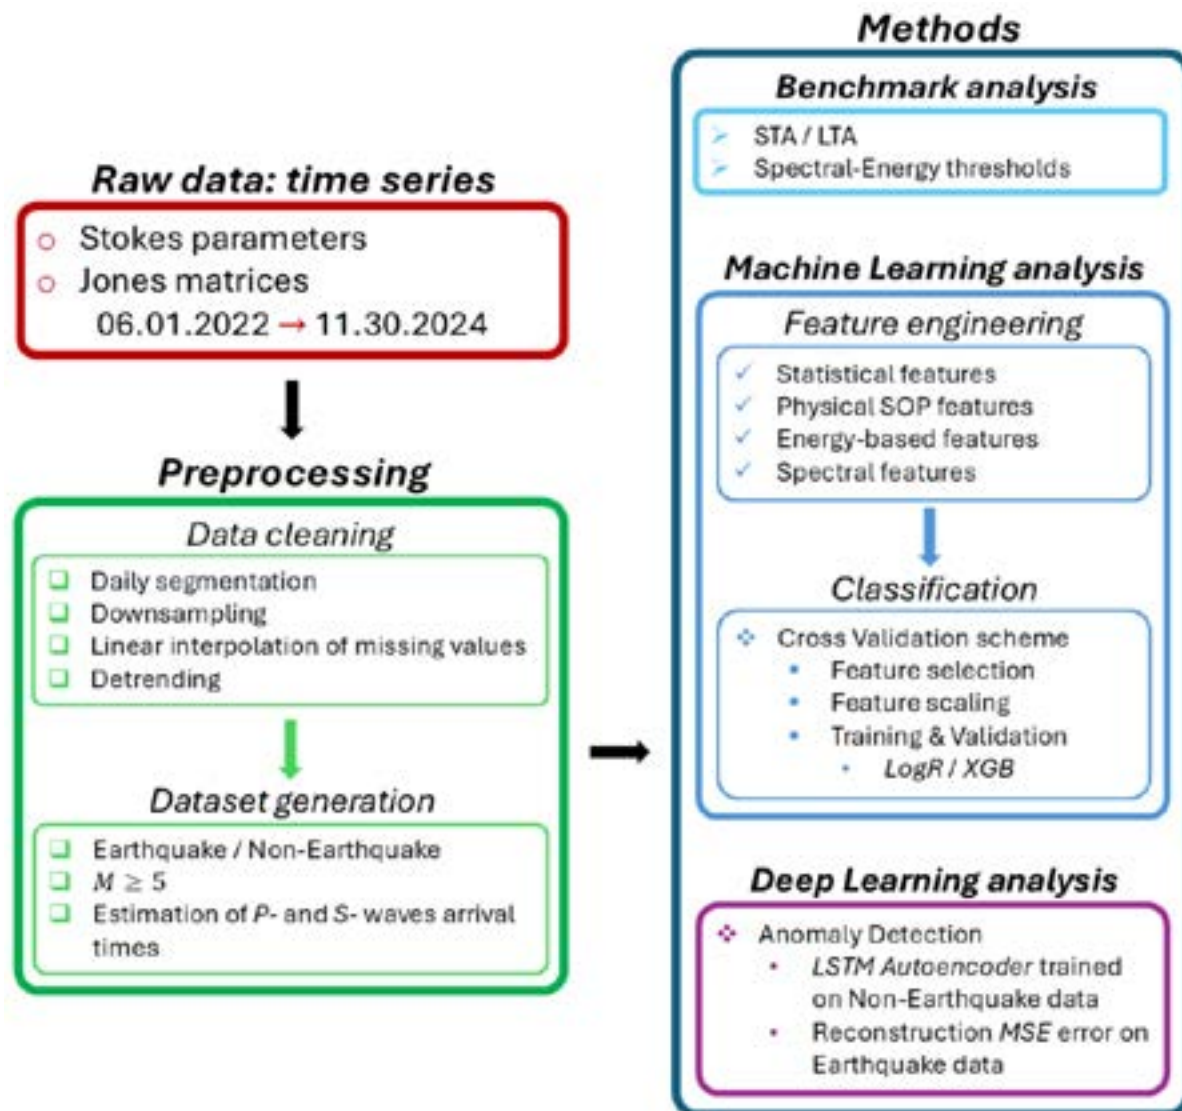
  

DL Model	RMSE	MAE	R <sup>2</sup>
Autoencoder	0.327 ± 0.175	0.259 ± 0.137	0.855 ± 0.278

**Table 2. Machine Learning and Deep Learning performances.**

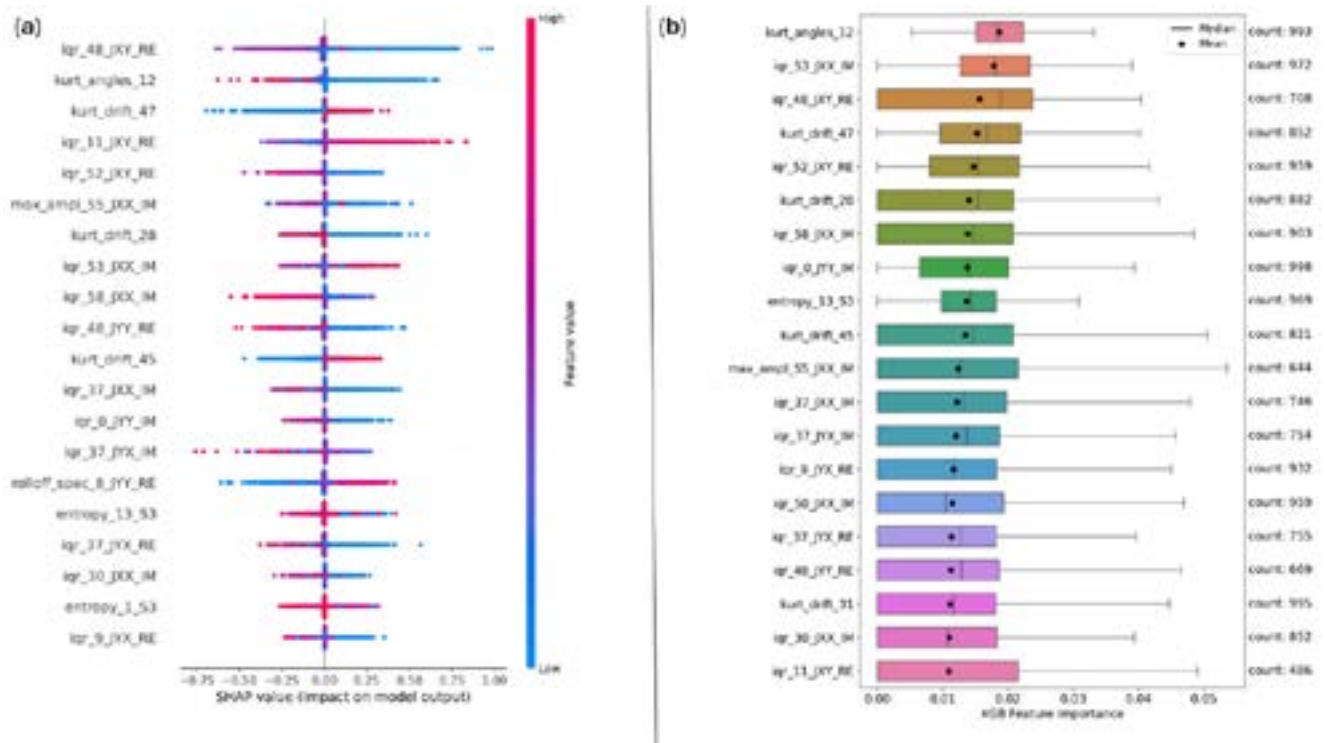
**Top:** Classification task. The classification pipeline was used to detect whether the SOP signal is affected by perturbations due to seismic events. XGBoost (XGB) and Logistic Regression (LogR) were trained and validated using repeated stratified cross-validation ( $k = 10$  folds,  $n = 100$  repetitions). Reported values represent median ± interquartile ranges. A t-test applied to the metric differences showed that the outputs of the two models are statistically different.

**Bottom:** Reconstruction task. The Autoencoder was evaluated using a train/validation split (approximately 250,000 training sequences and 25,000 validation sequences). RMSE, MAE, and  $R^2$  were computed on the validation set, and median ± interquartile ranges are reported.

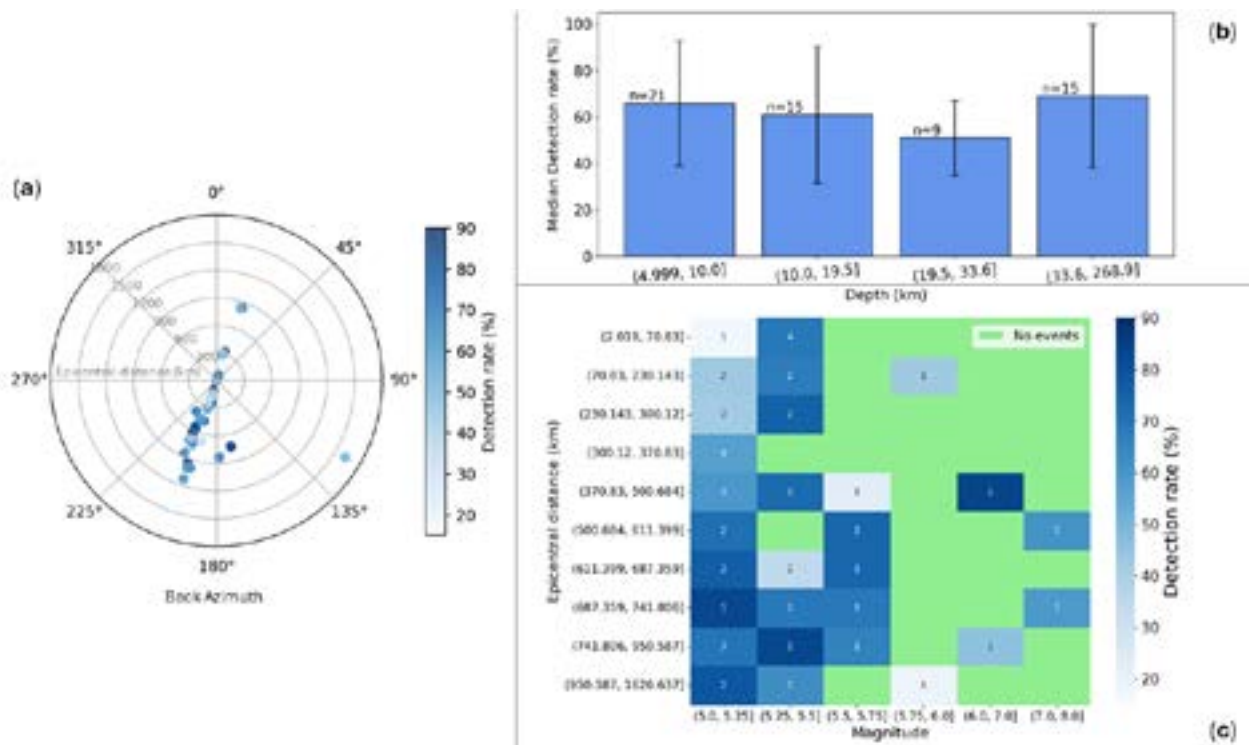


**Figure 1.**

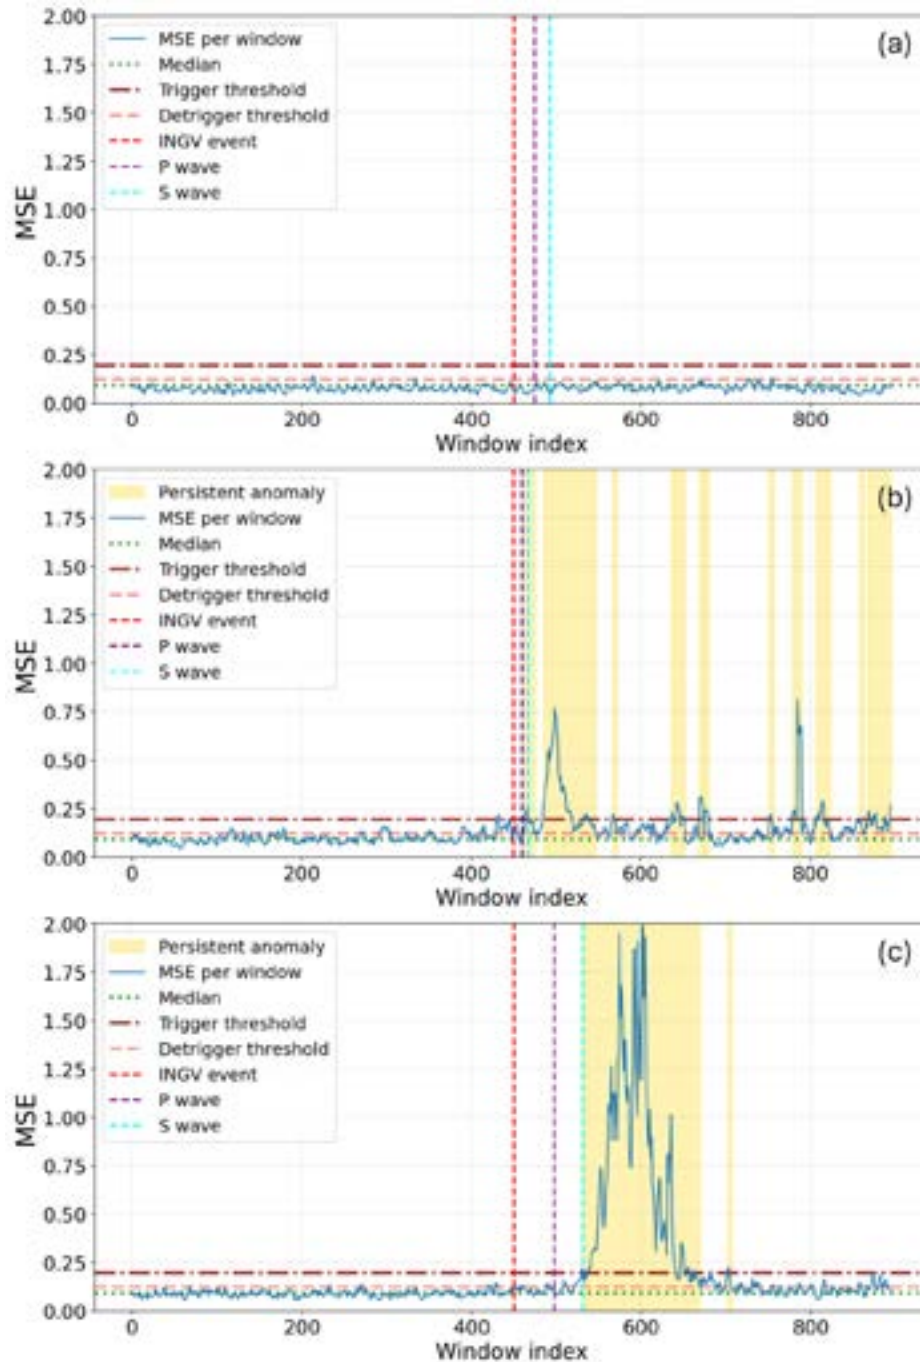
**Flowchart.** Overview of the data processing and analysis workflow. The raw dataset consisted of daily State of Polarization (SOP) time series (Stokes parameters and Jones matrix components) recorded along the Med-Nautilus submarine optical cable, between June 2022 and November 2024. After data cleaning (daily segmentation, downsampling, interpolation of missing samples, and detrending), the 24h time series were structured into Earthquake and Non-Earthquake days, with seismic labels assigned using the Istituto Nazionale Geofisica e Vulcanologia (INGV) catalogue ( $M \geq 5$ ) and theoretical P- and S-wave arrival times estimated from a global velocity model, IASP91. Two baseline detectors were first evaluated: (i) a classical STA/LTA (Short Time Average / Long Time Average) amplitude-based trigger, and (ii) a fixed spectral-energy threshold approach. The Machine Learning pipeline built a feature space from statistical descriptors, physically motivated SOP features, energy-based measures, and spectral attributes. Then, feature selection, feature scaling, and repeated cross-validation to train Logistic Regression (*LogR*) and Extreme Gradient Boosting (*XGB*) algorithms were performed. A complementary Deep Learning analysis used an LSTM (Long-Short Term Memory) Autoencoder trained exclusively on Non-Earthquake days to identify seismic perturbations through reconstruction error. This integrated workflow provides a systematic framework to assess the detectability of seismic-induced polarization fluctuations using both classical and data-driven approaches.



**Figure 2. Feature importance analysis.** (a) SHapley Additive exPlanations (SHAP) summary plot showing the distribution of SHAP values for the 20 most influential features, averaged over the stratified and repeated cross-validation procedure ( $k = 10$ ,  $n = 100$ ). Each point represents the contribution of a feature to the model output for a single sample, with color encoding the feature value (blue: low, red: high). Positive (negative) SHAP values indicate an increased (decreased) predicted probability of the seismic class. (b) Ranking of the most influential features according to the internal Extreme Gradient Boosting (XGBoost) feature-importance measure, computed as average gain. Boxplots summarize the distribution of importance values across all cross-validation splits ( $counts = n \times k = 1000$ ), with black dots and vertical, continuous lines indicating mean and median values, respectively. Both approaches highlight a largely consistent set of relevant features, dominated by statistical descriptors related to robust measures of signal variability (e.g., interquartile ranges), with a smaller contribution from physically motivated features such as maximum amplitude, spectral roll-off, and entropy. The gradual decrease of importance values in the rankings suggests that the classifier relies on a heterogeneous set of comparably informative features rather than on a small number of dominant predictors.



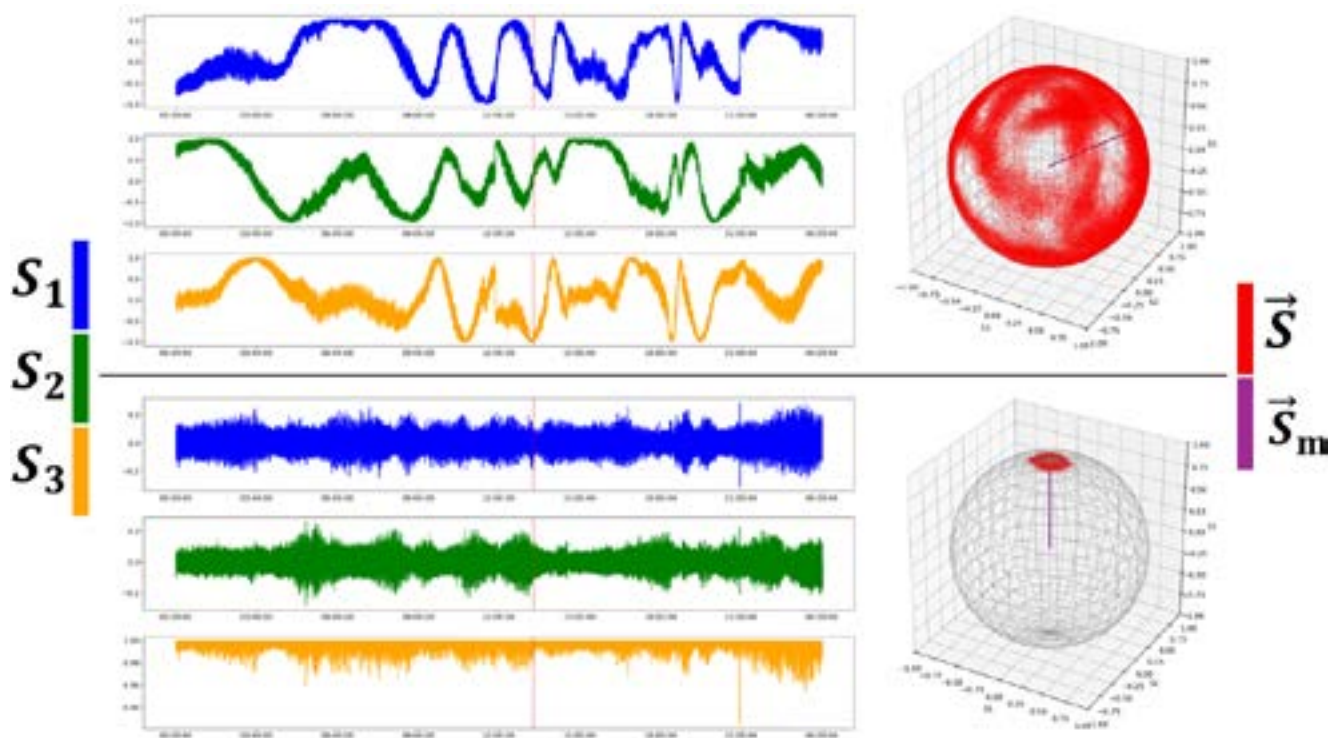
**Figure 3. Conditional detectability of seismic events based on Extreme Gradient Boosting performance.** The detection rate is computed for each of the 60 earthquakes with  $M \geq 5$  as the fraction of cross-validation repetitions in which the event is assigned to the seismic class by Extreme Gradient Boosting (XGBoost). **(a)** Detection rate as a function of back-azimuth and epicentral distance. Each point represents an earthquake, and is coloured according to its detection rate. Events cluster predominantly in the  $180^\circ$ – $225^\circ$  sector, consistently with the regional seismicity relative to the Med-Nautilus cable orientation. However, detection rates vary substantially even for similar back-azimuths and distances, indicating the absence of a clear angular or geometric control on detectability. **(b)** Median detection rates and error bars, computed focusing on events within each quartile of the depth distribution. The comparability of values across depth classes suggests the absence of a systematic dependence of detectability on depth. **(c)** Detection matrix as a function of magnitude and epicentral distance (10 quantile bins for distance). Numbers inside cells indicate the counts of events in each bin, and colours represents their detection rate; green cells correspond to bins for which no seismic events occurred in the observation period. No consistent trend of detection rate with magnitude or distance emerges, partly due to the limited number of events at higher magnitudes. Overall, the figure shows that detectability does not follow simple monotonic patterns in any single seismological variable, reflecting relevant event-to-event variability.




**Figure 4. Testing Autoencoder.** Earthquakes used to test our anomaly detection system. Vertical lines indicate key time references: the actual earthquake occurrence time given by Istituto Nazionale Geofisica e Vulcanologia (INGV) (in red), the estimated arrival time of P-waves (in violet), and the estimated arrival time of S-waves (in cyan). Horizontal lines represent the median reconstruction error from the training set (in green), the trigger threshold  $T$  (in brown) and the dettrigger threshold  $DT$  (in salmon). The yellow bands in the background indicate intervals where the system detects persistent anomalies according to the trigger, de-trigger, and temporal duration parameters. Panel (a) corresponds to an earthquake of magnitude  $M = 5.5$ , the middle plot panel (b) to an earthquake of  $M = 5.8$  and panel (c) refers to an earthquake of  $M = 7.9$ . In the first plot, the system does not detect any notable anomalies in the State of Polarization (SOP) signals ( $S_1$  and  $S_2$ ), despite the actual occurrence of a seismic event. For the second and third plots, the system successfully identifies anomalies in the SOP signals that we interpret as perturbations induced by seismic events. This behavior is particularly evident in panel (c), where the reconstruction error is very high.




**Figure 5. The Med-Nautilus section of interest.** Med-Nautilus comprises 5,729 km of cable, with a design capacity of 3.84 Tb/s distributed across six fiber optic pairs. In this work, we focused on the Catania (ITA) - Tel Aviv (ISR) section. This cable section is represented by the dashed red line: the actual position of the cable may slightly differ from this representation, since cable precise coordinates are not publicly available due to security reasons. The cable section was approximated by the geodesic  $CT$  (green continuous line); this section has length above 2,500 km. The circles shown in the figure represent the 60 seismic events with magnitude  $M \geq 5$  that occurred in the Mediterranean area between June 1<sup>st</sup>, 2022 and November 30<sup>th</sup>, 2024. The aim of this research work is studying the changes in the State of Polarization (SOP) of the electromagnetic signal in the Med-Nautilus  $CT$  fiber cable section due to the mechanical stress derived from P- and S- seismic waves. The two geodesics passing through  $C$  and  $T$  (green dashed lines) create a sort of wedge-shaped region which is our zone of interest; the outer region is denoted as  $S$  zone.



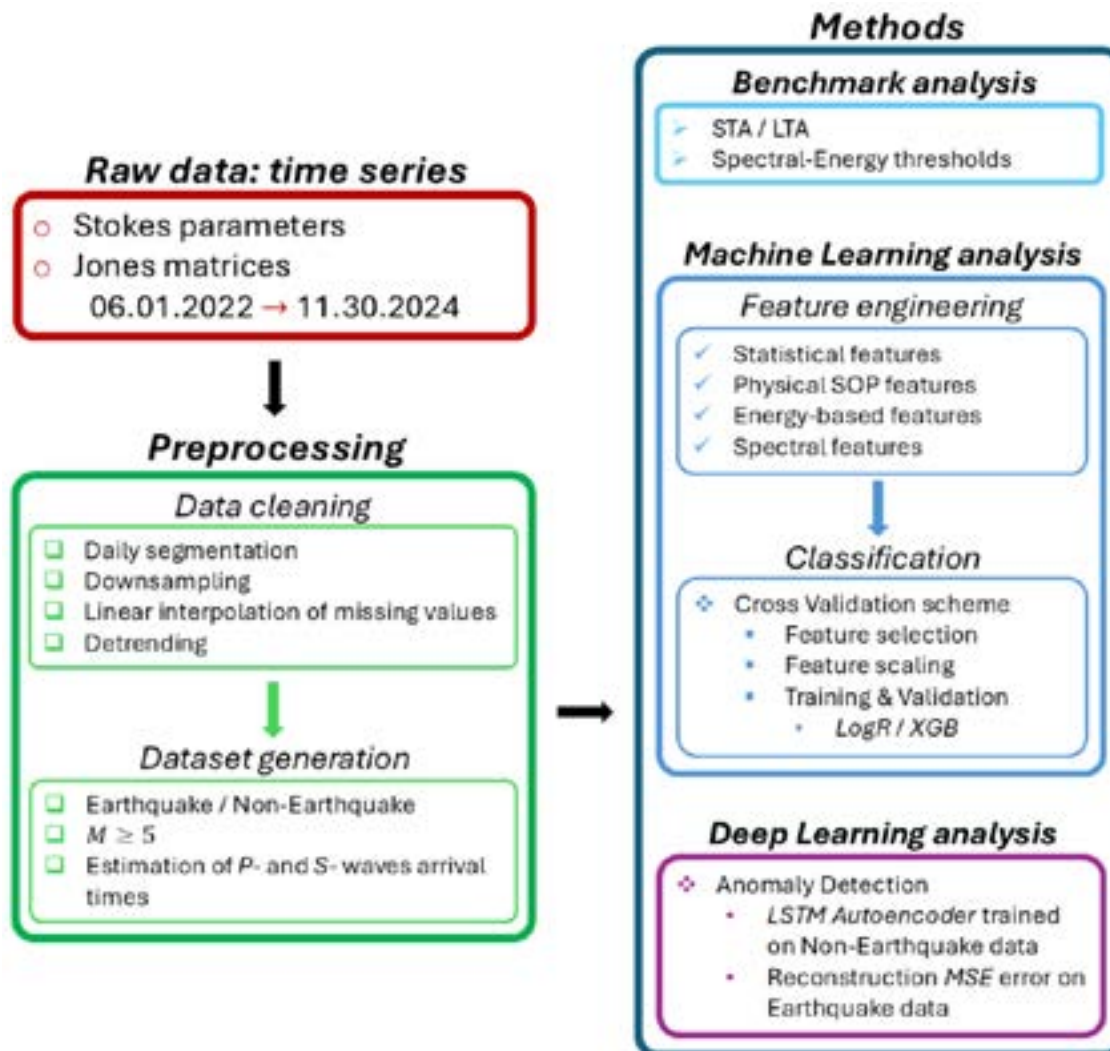
**Figure 6. Stokes parameters: before and after preprocessing.** Representation of the three Stokes parameters on 14<sup>th</sup>, February, 2023. Left panels report the time series corresponding to the three parameters, with a vertical red line indicating the time, provided by Istituto Nazionale Geofisica e Vulcanologia (INGV), of a seismic event of magnitude  $M = 5.8$ , which occurred on that day in Romania (Latitude = 45.0932, Longitude = 23.2446, Depth = 10 km). Right panels show the Poincaré spheres and the trajectory of the polarization state vector  $\vec{S}$  on their surface throughout the day. The top panels display the signal before de-trending, while the bottom panels show the de-trended signal. De-trending is performed by computing the mean Stokes vector  $\vec{S}_m$  every 2s, determining the angle  $\theta_m$  between  $\vec{S}_m$  and the z-axis (aligned with the third Stokes parameter,  $S_3$ ) and rotating all vectors  $\vec{S}$  in that window by  $\theta_m$ . This operation aligns the third Stokes  $S_3$  parameter to one. Since the sum of the squares of the three Stokes parameters is equal to one, the number of independent parameters is effectively reduced to two.





**Figure 7. Autoencoder input: tensor visual representation.** In this figure, there is an illustration regarding the single tensor  $S_i$ ,  $i \in \{1, \dots, 280\}$ . Each tensor contains the two independent Stokes parameters  $S_{i,1}$  and  $S_{i,2}$ , each  $\Delta t = 30 \text{ min}$  long. Each time series  $S_{i,c}$ ,  $c \in \{1, 2\}$ , is divided into shorter subsequences  $s_{ij,c}$ . These sub-sequences are  $\delta t_{sub} = 10 \text{ s}$  long (purple arrows). A new sub-sequence is generated by moving forward in time a window of  $\delta t_{sub}$  with  $\delta t_{stride} = 2 \text{ s}$  stride (green arrows), resulting in an overlap of  $\delta t_{overlap} = 8 \text{ s}$  between consecutive segments (light blue arrows). Given this temporal segmentation, the total number of subsequences  $s_{ij,c}$  extracted from each  $S_{i,c}$  is  $N = \left\lfloor \frac{\Delta t - \delta t_{sub}}{\delta t_{stride}} \right\rfloor + 1 = 896$ .

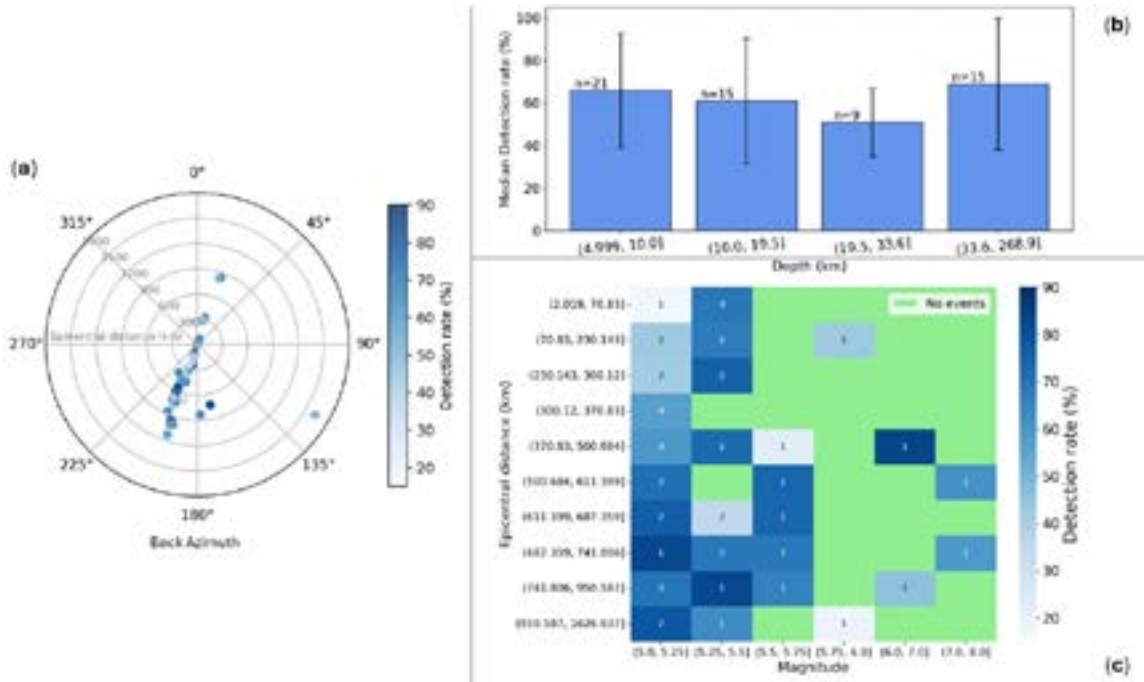


**Figure 8. Training reconstruction error distribution.** Training reconstruction error distribution given by the Autoencoder. Each  $MSE_{ij}$  is evaluated on a single sub-sequence  $s_{ij,c}$ . The anomaly threshold was set equals to the upper inner fence of the distribution, defined as  $\eta = Q_3 + \frac{3}{2} \cdot IQR$ .

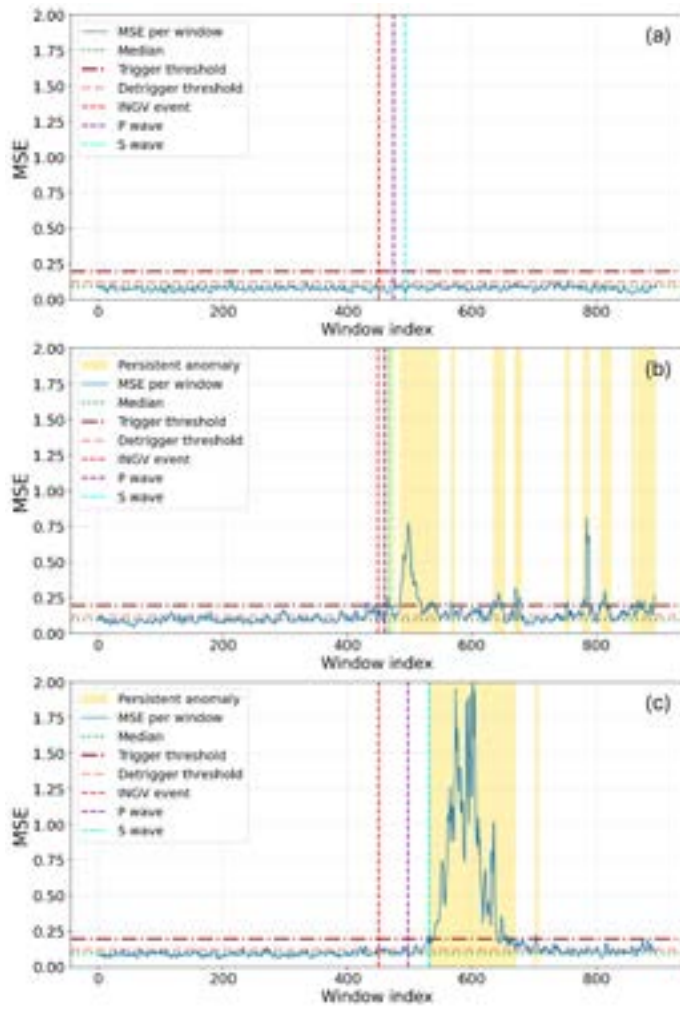




ARTIC

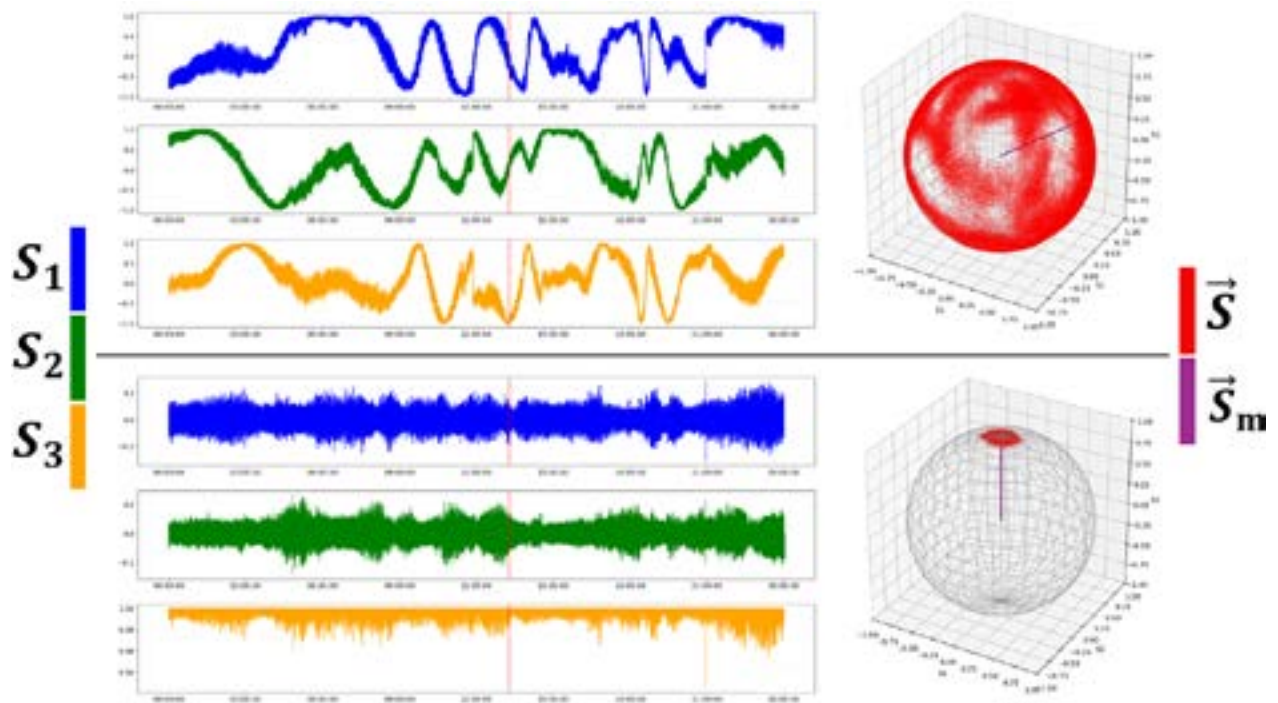



ARTI






ARTICLE





ARTIC



ART