



Unsupervised, Modulation-Independent Detection of Harmful Fiber Events from SOP Signatures in a Live Metro Ring

Downloaded from: <https://research.chalmers.se>, 2026-10-24 07:18 UTC

Citation for the original published paper (version of record):

Sadighi, L., Furdek Prekratic, M., Karlsson, S. et al (2027). Unsupervised, Modulation-Independent Detection of Harmful Fiber Events from SOP Signatures in a Live Metro Ring. Optical Fiber Communication Conference (OFC 2027)

N.B. When citing this work, cite the original published paper.

Unsupervised, Modulation-Independent Detection of Harmful Fiber Events from SOP Signatures in a Live Metro Ring

Leyla Sadighi^{1*}, Stefan Karlsson², Devika Dass¹, Marija Furdek³, and Marco Ruffini¹

¹Department of Computer Science and Statistics, IRIS Research Group, ADAPT Centre, Trinity College Dublin (TCD), Dublin, Ireland

²Micropol Fiberoptics AB, Halmstad, Sweden

³Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

*sadighil@tcd.ie

Abstract: We propose a stand-alone, unsupervised SOP-monitoring framework, requiring no labels or hyper-parameter tuning, with channel alignment to mitigate modulation-format dependence. On a live metro ring, it achieves 97–99% accuracy for CW/DP-16QAM and 79–88% for QPSK.

1 Introduction

Optical fiber networks are continuously exposed to various physical-layer events, ranging from routine handling and bending to vibrations that can precede a fiber cut, or eavesdropping that extracts light through a bend. Operators must detect these events promptly to protect service availability and data confidentiality. Among the various monitoring techniques, tracking the changes in the state of polarization (SOP) of transmitted light with a low-cost polarization analyzer is a viable way of turning any fiber into a sensor [1]. Supervised machine learning (ML) models excel in event classification based on SOP signatures of different events [2–4]. However, these techniques require staging of every observable event on the fiber, recording the data and annotating it by an expert before a detector can be trained to learn the events’ signatures, which is rarely feasible for a network operator. A practical detector must, therefore, be able to detect potential threats without prior knowledge of the event type. Unsupervised learning (USL) approaches meet this requirement, as they learn the structure of the monitored data directly and expose anomalous samples as outliers, typically through clustering. Their role in SOP-based anomaly detection has so far been explored only on laboratory data with unmodulated, continuous-wave (CW) signals with label-based hyperparameter tuning [1, 5]. In deployed networks, the modulation format (MF) of the monitored channel changes the level and noise floor of the SOP signature, and field noise adds fluctuations unrelated to any event, so both can hide the events from a USL learner.

In this paper, we propose a novel, fully unsupervised framework for SOP-based monitoring that does not require any ground-truth event labels neither for model training nor for hyper-parameter selection. We develop a dedicated 3-step data pre-processing chain which ensures that the events, rather than noise or the MF, dominate the USL model. The key part of the chain is the channel alignment step that standardizes every monitored channel with its own statistics, allowing a single model to detect events across the different MFs without requiring their prior separation.

We experimentally validate our framework in the noisy, live production metro ring of Asiera, Ireland’s national research and education network operator in Dublin. The network is exposed to eight physical-layer events, ranging from routine switching and bending to mechanical vibrations and eavesdropping. We record SOP signatures for unmodulated CW light, a 200G dual-polarization 16-ary quadrature amplitude modulation (DP-16QAM) channel, and a 100G quadrature phase-shift keying (QPSK) channel. We evaluate 13 USL approaches and benchmark them against the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) method of [5], and the supervised learning (SL) approach from [3] trained on labeled data. Events are detected with 97–99% mean accuracy for CW and DP-16QAM signal, approaching the SL baseline. Our channel alignment mechanism fully compensates the initial QPSK performance degradation in both USL and SL, enabling a single USL model to operate across different MFs.

2 Real-World Experimental Setup and Collected Signatures

Fig. 1 shows the experimental setup of the OpenIreland testbed at TCD and the 125 km Dublin metro ring of Asiera. We launch light from one source at a time: an unmodulated CW laser, a 200G DP-16QAM or a 100G QPSK channel from a commercial coherent transponder. The light traverses six Reconfigurable Optical Add-Drop Multiplexer (ROADM) nodes of the ring in a 400 GHz spectrum-as-a-service window and returns to the laboratory. There, the events are applied on a fiber patch under test, and the received light feeds the Micropol optical analyzer (OA). Its polarization beam splitter (PBS) separates the horizontal (H) and vertical (V) polarization components, which two photodiodes (PD-H and PD-V) convert into the intensities I_1 and I_2 , sampled by an analog-to-digital converter (ADC) at 44 kSa/s. The OA computes $\Delta S = I_1 - I_2$ and outputs the numerical polarization state variation (NPSV), $\text{NPSV}_t = \Delta S_t - \Delta S_{t-1}$, which we record. The relaxed fiber (no-event (NE)) is the normal, baseline state. Four events model non-harmful operation: 1) fast switching (FS) and 2) slow switching (SS), where a motorized polarization controller changes only the SOP of the light, abruptly in FS and gradually in SS, 3) macro-bending (MB) as routine handling with soft bending, and 4) tapping (TAP), which produces short taps on the fiber that cause brief SOP transients.




Fig. 1. Experimental setup: one of three sources is launched from the OpenIreland laboratory into the Asiera metro ring and returns after 125 km to a fiber patch on which the eight events are emulated.

Four events model harmful operation: 5) 80 Hz vibration (VB), 6) two-tone 80+135 Hz vibration (VB2), and 7) 80 Hz acoustic excitation by a loudspeaker (VB-LS) include 80 Hz vibration as a possible cut predecessor [5], while 8) eavesdropping via bending (EAV) emulates eavesdropping using a commercial clip-on coupler. Each event is recorded for 296 s, and a 4096-point fast Fourier transform (FFT) of every second of the NPSV yields one data point of 2048 unique frequency bins, i.e., half of the 4096 bins up to the Nyquist frequency. The 296 data points of one event and one MF form its SOP signature, a matrix of 296 data points by 2048 bins, as shown in Fig. 1.

3 Unsupervised Detection Framework

The proposed USL pipeline performs the task of detecting the *event* vs. *no-event* status across eight events and three MFs (24 tasks in total), each fed as a monitoring stream of 592 data points (296 for the *event* and 296 for the *no-event* state). To this end, we develop three label-free preprocessing steps as front end, followed by clustering by the learners, as illustrated in Fig. 1 and described below.

Front-end. We design three pre-processing steps to remove noise before clustering, which is typically the most differing factor among data points, misleading the clustering learners. First, *log-frequency pooling* merges neighboring frequency bins into 76 bands, keeping fine detail at the low frequencies where events such as vibrations appear. This averages out noise while preserving the event. Second, *temporal context aggregation* replaces each data point by the average of the last W data points, smoothening brief fluctuations while keeping the longer-lasting events. Finally, *channel alignment* rescales each monitored channel by its own average and spread. When channels of different MFs are monitored together, their sharing of common scale allows the learner to separate events rather than MFs.

Learners. A clustering learner splits the processed data points into two groups corresponding to *event* vs. *no-event*. We tested 13 USL approaches and focus only on the top three: Uniform manifold approximation and projection (UMAP)-hierarchical density-based spatial clustering of applications with noise (HDBSCAN) places similar data points close together on a simplified map and then groups the points that gather in the same area. Non-negative matrix factorization (NMF)-Gaussian mixture model (GMM) summarizes each data point by a few features and models each group as a bell-shaped cloud. Self-tuning (ST)-Spectral connects similar data points and cuts the weakest connections.

Choosing the settings, such as W , by accuracy would require the ground truth. Instead, our framework selects the settings which maximize a *quality index* $Q = s \times c$ computed from the data alone. The silhouette s is high when the two event data groups found by the clustering model, ideally *no-event* and *event*, are compact and well separated, and the temporal coherence c is high when consecutive data points stay in the same group, as fiber states persist in time.

4 Results

The performance of the top three approaches, i.e., UMAP-HDBSCAN, NMF-GMM and ST-Spectral is reported in Fig. 2 in terms of the mean accuracy over all 32 tasks. We compare them with the DBSCAN detector of [5], applied to the data without our front-end and with its settings chosen by the same quality index, and with a SL reference, i.e., histogram gradient boosting (HGB) trained with the full ground truth on the same spectral data under a stratified 80/20 hold-out with 5-fold cross-validation. With CW and DP-16QAM signals (Fig. 2(a),(b)), the three USL approaches reach a mean accuracy over the eight events of 99.3–99.5% and 97.2–99.5%, respectively. This matches the SL reference at 99.2% and 99.7% and far exceeds the DBSCAN baseline, which reaches only 67.0% with CW. For the QPSK signal (Fig. 2(c)), the mean accuracy of the three USL approaches drops to 79.1–88.1%, and that of the




Fig. 2. Accuracy (%) of the three best unsupervised approaches, the DBSCAN baseline [5] and the supervised reference on the eight binary detection tasks (no event vs. event).




Fig. 3. (a) Accuracy (%) of the three best approaches with the three MFs pooled naively or channel-aligned. (b), (c) Principal components of the pooled NE and EAV data points without and with channel alignment. (d) Mean accuracy of the three best approaches vs. temporal context W , band = range, $W > 8$ with QPSK only.

SL reference to 91.3%. Fig. 2(d) shows one model monitoring all three MFs at once, as an operator with many wavelengths would need. Our proposed channel alignment step enables the single USL model to detect the events with 80.4–87.3% mean accuracy, against 96.8% for the SL and 73.2% for the DBSCAN baseline, respectively. The benefits of channel alignment are further supported by Fig. 3. As Fig. 3(a) shows, naive pooling of the three MFs without channel alignment fails, and the three approaches fall to 52.3–55.9% mean accuracy, close to chance. Fig. 3(b) illustrates the reason using the EAV scenario, where the data points of both states are clustered per MF. With channel alignment, the MFs overlap and the two states are separated, as shown in Fig. 3(c). Fig. 3(d) further analyses the importance of the temporal context, modeled by tuning the value of W . DP-16QAM reaches 94.6% already at $W = 1$, whereas CW requires $W \geq 4$ to rise from 78.5–85.1% to 97.7%. For QPSK, extending W up to 32 raises the accuracy of UMAP-HDBSCAN from 80.1% to 85.2%. Together, these results show that the proposed pre-processing chain enables clustering as accurate event detector that does not require ground truth, and that channel alignment allows a single model to monitor several MFs at once.

5 Conclusion

We propose the first fully unsupervised framework for SOP-based monitoring that does not require any ground-truth event labels neither for model training nor for hyper-parameter selection. A dedicated 3-step data pre-processing with channel alignment that allows a single model to detect events across the different MFs without requiring their prior separation. Eight physical-layer events are detected from SOP signatures of a deployed metro ring carrying customer traffic, achieving 97–99% mean accuracy for CW and DP-16QAM signals, matching the supervised learning baseline.

Funding and Acknowledgments

This work was supported by the CELTIC-NEXT SUSTAINET-Advance project through Enterprise Ireland (IR-2025-0075) and VINNOVA (2025-02987), Research Ireland under Grants 18/RI/5721 (OpenIreland Research Infrastructure) and 23/RC/13506 (RINN AI), and VR (2023-05249). The authors gratefully acknowledge Asiera for access to its deployed metro network.

References

1. S. Pellegrini et al. *J. Opt. Commun. Netw.*, 17(2):A196–A209, 2025.
2. L. Sadighi et al. *J. Lightwave Technol.*, 43(13):6092–6101, 2025.
3. L. Sadighi et al. *J. Lightwave Technol.*, 44(16):7274–7283, 2026.
4. K. Abdelli et al. *J. Lightwave Technol.*, 43(4):1902–1914, 2025.
5. L. Sadighi et al. *IEEE Trans. Netw. Serv. Manag.*, 23:432–442, 2026.