CHALMERS
UNIVERSITY OF TECHNOLOGY

Towards Generalizable ML-Based Event Recognition Using Φ-OTDR and Siamese Networks

Downloaded from: https://research.chalmers.se, 2026-05-13 22:23 UTC

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

Nogueira Ribeiro, A., Lobato, F., Costa, J. et al (2026). Towards Generalizable ML-Based Event Recognition Using Φ-OTDR and Siamese Networks. Proceedings of the 30th International Conference on Optical Network Design and Modelling (ONDM)

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

research.chalmers.se offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all kind of research output: articles, dissertations, conference papers, reports etc. since 2004. research.chalmers.se is administrated and maintained by Chalmers Library

(article starts on next page)

Towards Generalizable ML-Based Event Recognition Using Φ-OTDR and Siamese Networks

Andrei N. Ribeiro *, Fabrício R. L. Lobato *, João C. W. A. Costa *, Paolo Monti †, Carlos Natalino †

* Institute of Technology, Federal University of Pará, 66075-110 Belém, Pará, Brazil
andrei.ribeiro@itec.ufpa.br; {frl, jweyl}@ufpa.br
† Department of Electrical Engineering, Chalmers University of Technology, 412 96 Gothenburg, Sweden
{mpaolo, carlos.natalino}@chalmers.se

***Abstract*—Machine learning models for event recognition (detection and classification) in optical fiber sensing often fail to generalize across deployments and require extensive retraining for new event types. This limitation poses challenges for practical deployment, particularly when novel event types emerge and system configurations change frequently. We propose an attention-weighted multi-similarity Siamese neural network (MS-SNN) for few-shot event recognition in distributed acoustic sensing applications. By combining five complementary similarity metrics with class-balanced episodic training, our approach learns generalizable embeddings from limited labeled data. The architecture enables both classification of known event types and detection of novel event types without model retraining. The method was trained on 5 out of the 9 classes available in the dataset. Then, evaluated on the entire 9-class dataset, our method achieves 97% accuracy for binary event detection with 98% recall using only 5-10 support samples per class. Our results also indicate that standard accuracy metrics mask performance disparities on imbalanced data, and that balanced accuracy provides a clearer understanding of model performance. We release an open-source implementation to facilitate reproducibility and accelerate research in generalizable optical network sensing.**

***Index Terms*—Distributed acoustic sensing, Siamese neural network, few-shot learning, anomaly detection, fiber sensing.**

I. Introduction

Optical fiber sensing technologies have emerged as powerful tools for infrastructure monitoring and perimeter protection in optical networks [1]. Multiple sensing modalities exist: optical time-domain reflectometry (OTDR) characterizes fiber faults and anomalies; distributed acoustic sensing (DAS), also known as Φ-OTDR, detects acoustic vibrations along the fiber with meter-scale spatial resolution [2]; and state of polarization (SOP) monitoring identifies mechanical disturbances and eavesdropping attempts using coherent transceivers [3]. These technologies share common signal processing pipelines: high-dimensional time-series data transformed to the frequency domain via fast Fourier transform (FFT), followed by machine learning (ML)-based recognition, and deployment across diverse environmental conditions [4].

ML approaches, particularly convolutional neural networks (CNNs) and multi-layer perceptrons (MLPs), achieve high accuracy for event classification within controlled settings, reaching up to 98.6% for SOP-based detection [3], 99.7% for XGBoost-based classification [5], and 91% for DAS-based recognition [6]. However, evidence shows that these supervised classifiers may suffer from critical generalization failures. Models trained on one fiber link or spectral band fail dramatically when deployed on different systems, with accuracy dropping from 98.6% to as low as 8.1% in cross-system scenarios [7]. Real-world deployments face additional challenges: the emergence of novel event types not seen during training, severe class imbalance in the training data, and the need for costly retraining when new events or system configurations are encountered [8], [9]. These limitations apply across sensing modalities, as OTDR, Φ-OTDR, and SOP systems all exhibit similar generalization gaps when operating under the *closed-world* assumption [7], [10], [11]. The closed-world assumption requires that all event types present during deployment must be represented in the training data, which is often impractical for optical network environments. An open-world setting, by contrast, allows the system to encounter event types absent from training and to flag them as novel rather than force-fit them into known classes. Siamese neural networks (SNNs) represent promising alternatives to mitigate these challenges by learning transferable similarity metrics rather than fixed class boundaries [12]–[15].

We propose an MS-SNN designed for generalizable event recognition across optical fiber sensing applications. The MS-SNN learns an embedding space in which similar events cluster, enabling classification from few labeled examples and detection of novel event types dissimilar to any known classes. The architecture integrates five complementary similarity met-

Paper accepted for presentation at ONDM 2026.
This work has been supported in part by the Horizon Europe ECO-eNET project with funding from the SNS JU under grant agreement no. 101139133, by the Swedish Research Council (grant agreement no. 2022-04798), and by the National Council for Scientific and Technological Development (CNPq), Brazil. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through ga no. 2022-06725.

The full open-source implementation of the models in this work is available at https://github.com/carlosnatalino/ONDM_2026_Siamese/.

rics: L1, L2, cosine, element-wise product, and a learned attention-weighted fusion. It is trained with class-balanced episodic sampling to handle the severe imbalance common in sensing and failure-management datasets.

Our contributions are threefold. First, we introduce an MS-SNN architecture for fiber-sensing event recognition and evaluate ir on a publicly available DAS dataset [2], comparing against supervised baselines and demonstrating both few-shot classification and novelty detection. Second, we show that classification accuracy alone is insufficient on imbalanced sensing datasets: balanced accuracy and confusion matrices reveal that high accuracy can mask poor performance on minority anomaly classes. Third, we release an open-source implementation to facilitate reproducibility and accelerate research in generalizable optical network sensing. Beyond this single Φ-OTDR benchmark, the modality-agnostic formulation offers a path toward few-shot, open-world recognition across optical-network sensing and failure management.

II. Related Work

Multiple fiber sensing technologies enable event detection and classification [16], [17]. OTDR characterizes fiber faults via backscattered pulses, providing fault localization along the fiber length. Φ-OTDR detects acoustic vibrations via Rayleigh backscattering, enabling perimeter protection with meter-scale resolution [6], [10], [18]. SOP monitoring exploits polarization variations in coherent transceivers to detect mechanical disturbances without dedicated equipment. All modalities share a common pipeline: acquisition, frequency-domain transformation via FFT, and classification using CNNs or other discriminative models. Supervised baselines reach up to 91.4% on public DAS datasets [2], 96.67% with lightweight architectures [18], and 98.57% for SOP-based detection [3].

These models achieve high accuracy within their training domain but exhibit critical generalization failures when deployed across different systems [7], [11]. For example, SOP-based classifiers trained on and O-band, 21 km link drop to 8.1% accuracy when tested on a C-band, 77 km system, versus 98.6% intra-system [7]. Surveys identify model generalization as a critical unresolved challenge [8], [9], [11]. Multi-system training improves generalization (91.1% accuracy) but requires labeled data from every target configuration [7]. Data-scarcity approaches include variational autoencoders with generative adversarial network frameworks reaching F1-score of 0.8 with 1% training data [19], and unsupervised anomaly detection [20]. A recent adaptive classifier generalizes to unseen attacks with 0.936 balanced accuracy [21].

Metric-learning alternatives avoid the closed-world assumption through pairwise comparison [12]. SNNs reach 92% accuracy on 20-way one-shot classification with 70.3% domain generalization to MNIST without retraining [12], and in optical networking applications it achieved 100% accuracy on unseen modulation formats [13], 83% accuracy on unseen fault categories [14], and 99.12% with zero false positives for performance-monitoring anomaly detection [15]. Few-shot theory highlights that reducing intra-class variation [22] and broadening training class diversity [23] matter more than support-set size. Prototypical networks build on this by representing classes as mean embeddings [24].

Closer to our problem, few-shot meta-learning has been applied to urban DAS monitoring using prototypical heads over power spectral density (PSD) features to classify novel third-party-intrusion events with 1–3 support samples [25], and cross-scene DAS generalization has been addressed through knowledge-distillation plus semi-supervised adaptation, reaching 90.6% accuracy with 50 target-scene labels [26]. Both approaches, however, still rely on a single scalar distance or a labeled source scene.

Our work differs on three axes. First, unlike prototypical or single-metric Siamese heads [13], [14], [25], the MS-SNN fuses five complementary similarity metrics via learned attention, avoiding over-reliance on any single geometric assumption. Second, unlike cross-scene adaptation approaches [26], the MS-SNN requires no labeled source scene or target-domain retraining, only a handful of support samples at inference. Third, we explicitly frame the open-world regime (novelty detection via distance threshold) and release an open-source implementation together with MLP and CNN baselines inspired by [6], addressing the reproducibility gap in current DAS and SOP literature.

III. Proposed Siamese Network Architecture

We now describe how the MS-SNN realizes the open-world recognition setting outlined above, from problem formulation through architecture and training to the evaluation protocols.

A. Problem Formulation

Given a set of frequency-domain representations $X = \{x_1, ..., x_N\}$ from optical fiber sensing systems with corresponding class labels $Y = \{y_1, ..., y_N\}$, we aim to learn an embedding function $f : \mathbb{R}^{d_{in}} \rightarrow \mathbb{R}^{d_{emb}}$ that maps signals to a compact embedding space. In this space, signals from the same event class should cluster together while signals from different classes should be separated, regardless of the specific sensing system or deployment configuration. The learned embedding enables prototype-based few-shot classification via nearest-neighbor search in the embedding space and novelty detection by measuring distances to known class prototypes. This approach contrasts with standard supervised classifiers, which require fixed output dimensions equal to the number of training classes. This formulation is agnostic to the sensing modality, meaning the same architecture should be applicable to OTDR, Φ-OTDR, or SOP data after appropriate frequency-domain transformation.

B. Multi-Similarity Siamese Network Architecture

A Siamese network is a pair of weight-sharing branches that embed two inputs and answer a single question: are they from the same class? This pairwise framing lets the same trained network be applied to event types unseen during training. Fig. 1 illustrates the proposed architecture. The embedding network follows the CNN architecture inspired by prior work

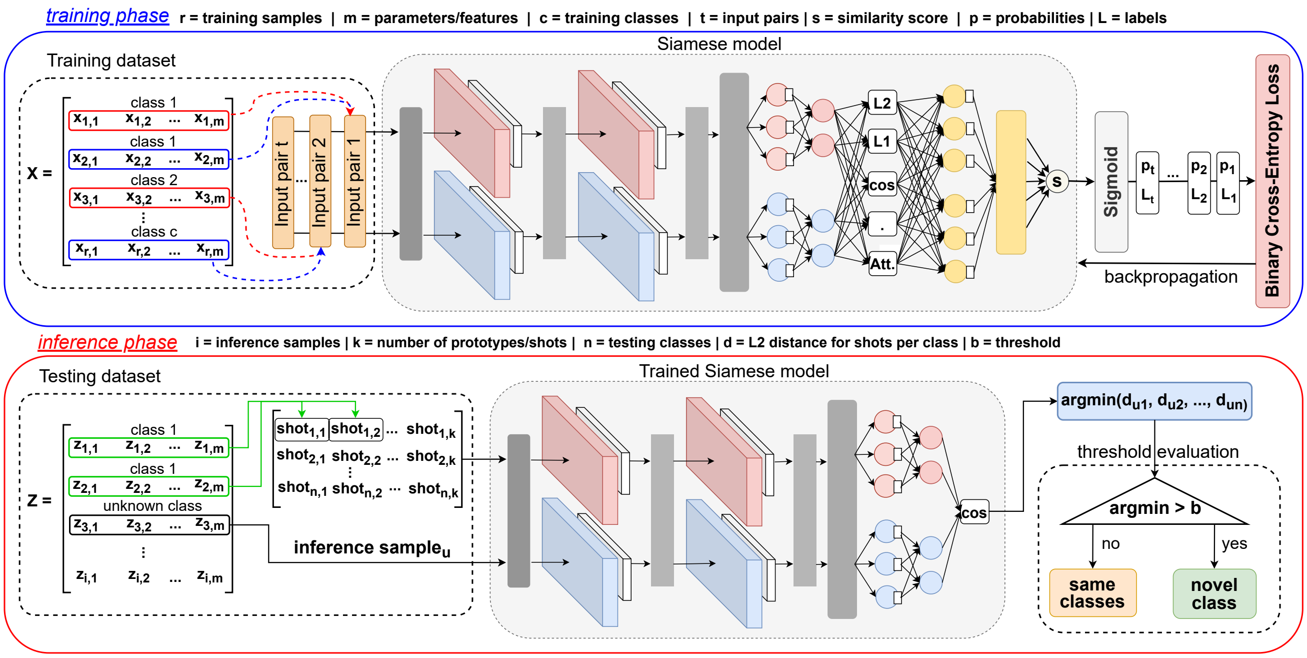

Fig. 1. Illustrative example of the training and inference phases of the proposed Siamese network architecture. During the training phase, pairs of samples are randomly drawn and assigned label 0 if the samples come from different classes, and label 1 if they come from the same class. During inference, a new sample is compared to k samples from each of the n known classes, either being classified as one of the classes or as a novel class (anomaly).

[6], adapted for Siamese weight sharing and multi-similarity. The architecture consists of an embedding block composed of convolutional layers with non-linear activations and max pooling, followed by fully connected layers that produce the embedding vector. The Siamese network is composed of two branches of the same embedding block, ensuring consistent feature extraction across paired inputs.

Between embedding pairs (e_1, e_2) we compute five complementary similarity scores, each capturing a distinct geometric property of the embedding space. L1 (Manhattan) is robust to outliers and sensitive to sparse, localized spectral differences. L2 (Euclidean) penalizes large deviations and serves as the standard proximity measure in metric learning. Cosine similarity is magnitude-invariant and captures directional alignment, useful when event signatures differ in shape but not in energy. The element-wise product highlights feature-level co-occurrence across embedding dimensions. Finally, a learned attention-weighted fusion combines the four scalar scores using a mechanism conditioned on the input embeddings, dynamically emphasizing the most informative metric for each pair. The resulting feature vector is processed by a comparison MLP to produce a similarity score in $[0, 1]$. This multi-metric design avoids over-reliance on any single distance measure, which matters given the heterogeneous acoustic signatures observed in real deployments.

C. Training Procedure

We employ class-balanced episodic sampling inspired by few-shot learning literature [24]: each episode samples pairs with a 50% positive (same class) and 50% negative ratio, with classes drawn uniformly regardless of dataset size so that minority classes receive adequate training signal. This prevents the model from collapsing onto majority-class shortcuts a common failure mode in imbalanced classification.

Training combines binary cross-entropy with multi-similarity regularization, $\mathcal{L}_{total} = \mathcal{L}_{BCE} + \lambda \cdot \mathcal{L}_{MS}$, where the regularizer penalizes small distances for negative pairs and large distances for positive pairs. Optimization uses Adam with weight decay and gradient clipping.

D. Evaluation Protocols

For few-shot classification, we compute class prototypes as mean embeddings of support samples. In an N-way K-shot evaluation, N classes are considered with K support samples per class. The model receives K labeled references per class and must classify new query samples from those alone via nearest prototype distance. For anomaly detection, samples are flagged as anomalous if their minimum distance to any prototype exceeds a threshold. This threshold is selected on the validation set by maximizing the F1-score for binary anomaly detection. We evaluate binary anomaly detection by grouping all event classes against the *regular* baseline class.

In practice, deployment proceeds in three steps: the operator collects a number of labeled support samples per known event class. Each set is summarized by its mean embedding, forming a class prototype. Every incoming query sample is then compared to the stored prototypes, being flagged as novel when its minimum distance exceeds the learned threshold.

IV. Performance Assessment

A. Experimental Setup

We evaluate our proposed method on the publicly available DAS dataset [2], which serves as a representative benchmark for fiber-sensing event recognition. The sensing system is a

TABLE I
MODEL HYPERPARAMETERS AND CONFIGURATION DETAILS.

Parameter	CNN	MLP	Siamese
batch size	256	256	128
learning rate	0.001	0.001	0.001
optimizer	Adam	Adam	Adam
epochs	100	100	100
early stopping patience	10	10	15
L2 regularization	0.0001	0.0001	0.0001
trainable parameters	∼1.2M	∼800K	∼1.5M

Φ-OTDR using an Optasense ODH-F interrogator operating at 1550 nm with a 20 kHz pulse repetition rate, 20 ns pulse width, and a 1 m spatial sampling interval over 1663 channels. The fiber is buried 1 m underground alongside a university campus pavement, capturing real-world urban events. The dataset contains 9 distinct classes: regular (baseline), car, walk, running, longboard, construction, fence climbing, fiber manipulation, and manhole open/close. Raw signals are processed using sliding windows (8192 samples, 2048 shift) and transformed to the frequency domain via FFT, retaining the first 2048 spectral components with Z-score normalization. This preprocessing pipeline follows standard practices in distributed fiber sensing literature, ensuring comparability with prior work.

We compare against two baselines: a CNN with Conv1D layers following the architecture from [6], and an MLP with fully-connected layers. Both baselines are trained on all 9 classes using standard supervised learning with an 80-10-10 train-validation-test split. The MS-SNN is instead trained on 5 of the 9 classes: regular, walk, car, fiber manipulation, and manhole open/close. These classes were chosen to pair the dominant baseline class with a diverse subset of event types spanning pedestrian activity, vehicular presence, and deliberate physical interactions. The remaining 4 classes are held out and used only at inference time to assess recognition of previously unseen events without retraining. Table I summarizes the hyperparameters. Models are evaluated on the test set using accuracy, balanced accuracy (which averages per-class recall so that a model is not credited for getting the dominant class right while ignoring the rare ones), F1-macro, and F1-weighted.

B. Classification Performance

Table II and Fig. 2 present the numerical results for all compared models. A critical observation is the substantial gap between accuracy and balanced accuracy for the supervised baselines. The CNN achieves 87.39% accuracy but only 63.78% balanced accuracy, revealing a gap of nearly 24 percentage points. Similarly, the MLP achieves 87.48% accuracy,

TABLE II
NUMERICAL RESULTS OF THE BENCHMARKED MODELS.

Metric	CNN	MLP	MS-SNN				
			9w-1s	9w-5s	9w-10s	9w-15s	9w-20s
Accuracy (%)	87.39	**87.48**	69.02	74.42	75.00	75.35	**75.90**
Balanced Acc. (%)	63.78	**69.67**	33.50	40.95	42.93	**44.13**	44.07
F1 Macro (%)	61.92	**67.54**	27.42	38.26	39.80	40.44	**41.69**
F1 Weighted (%)	85.63	**88.07**	69.60	75.09	75.65	75.82	**76.74**

Note: 9w = 9-way; 1s/5s/10s/15s/20s = 1-shot/5-shot/10-shot/15-shot/20-shot

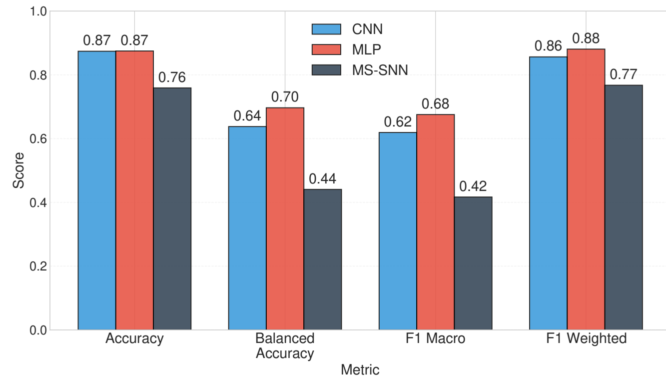

Fig. 2. Comparison among the CNN, MLP, and the proposed MS-SNN across multiple performance metrics: accuracy, balanced accuracy, F1-macro, and F1-weighted. MS-SNN results correspond to the best-performing configuration (9-way 20-shot).

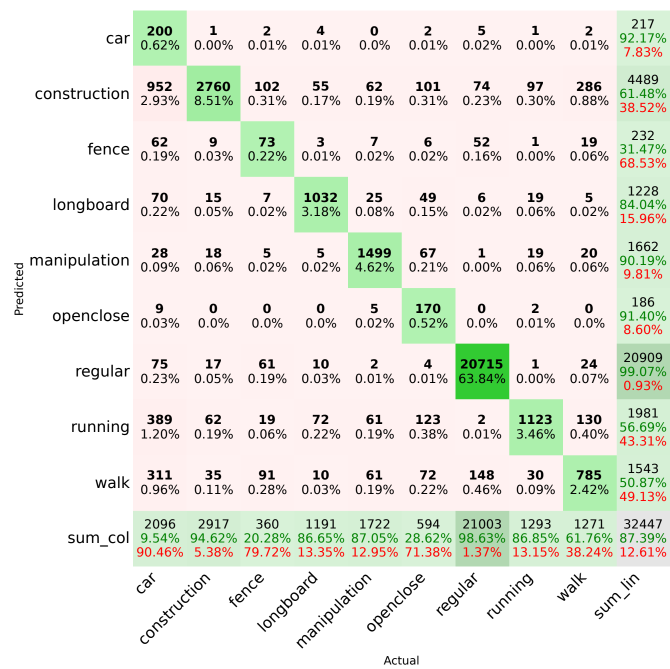

Fig. 3. Confusion matrix for the test set using the CNN with architecture inspired by [6]. Note the near-perfect recall on the dominant *regular* class (>98%) contrasted with poor recall on minority classes: *car* achieves only ∼10% recall with most samples misclassified as *construction*.

compared with 69.67% balanced accuracy. Interestingly, the MLP achieves higher performance than the CNN, while having 33% less trainable parameters (see Table I). This discrepancy indicates that both models perform well on the dominant *regular* class but struggle with minority anomaly classes, a critical issue for practical deployment.

The few-shot scaling analysis in Table II shows that MS-SNN performance improves from 69.02% (1-shot) to 75.90% (20-shot), with the largest gain between 1-shot and 5-shot and diminishing returns beyond 10 samples per class. Unlike adaptive classifiers that require incremental weight updates for each new labeled batch, our approach incorporates novel events immediately through prototype comparison alone.

The confusion matrices in Fig. 3 and Fig. 4 reveal class-specific patterns. The CNN achieves near-perfect recall on the *regular* class but poor recall on minority classes such as *car* and *fence*. The MLP is more balanced but still confuses between similar event types. A by-product of these matrices is that the baselines also perform implicit binary anomaly detection when predictions are grouped into regular vs. non-regular. The CNN's near-total recall on *regular* yields a very low false-alarm rate, but at the cost of missing several minority anomalies (e.g., *car*), which is exactly the failure mode that matters in perimeter protection.

The MS-SNN in the 9-way 20-shot configuration (Fig. 5) achieves 75.90% accuracy and 44.07% balanced accuracy, which is lower than the supervised baselines. This comparison must be read in context. The CNN and MLP are trained on thousands of labeled samples per class over the full training set, whereas the MS-SNN relies on only 20 support samples per class at inference. More critically, the supervised baselines operate under the closed-world assumption and cannot recognize event types absent from training without full retraining. The MS-SNN, by contrast, incorporates new event types immediately by providing reference samples, with no weight updates. This open-world capability is the primary motivation of the approach. Consistently with [22], the gap on fine-grained 9-class classification reflects intrinsic limits on small-backbone embedding capacity and high intra-class acoustic variability (e.g., different cars or walking styles produce markedly different spectra), compounded by pair-based episodic training providing weaker class-prototype supervision than N-way classification heads trained on thousands of samples. Still, Fig. 5 shows that episodic training yields more balanced per-class performance across minority classes.

C. Anomaly Detection Performance

For practical deployment, the primary task is detecting anomalies versus normal operation, as this is critical to preserving normal operating conditions. Framing the problem as binary classification (regular = normal, all other classes = anomaly), the MS-SNN achieves exceptional performance across all shot configurations. The model achieves approximately 97% accuracy and balanced accuracy, confirming that performance is not biased by class imbalance. F1-scores exceed 96% with precision around 94% and recall exceeding 98%. The high recall ensures that nearly all anomalies are detected, which is critical for perimeter protection and infrastructure monitoring applications.

D. Discussion

The results reveal a trade-off between fine-grained classification and anomaly detection. The MS-SNN's 44% balanced accuracy on 9-class classification is far below 97.5% binary anomaly detection, reflecting the inherent difficulty of separating similar anomaly types in frequency-domain representations. Rather than pushing a single model to cover both tasks, we argue for a hierarchical deployment: the MS-SNN acts as a high-recall open-world anomaly detector, and flagged windows are forwarded to a specialized classifier for categorization.

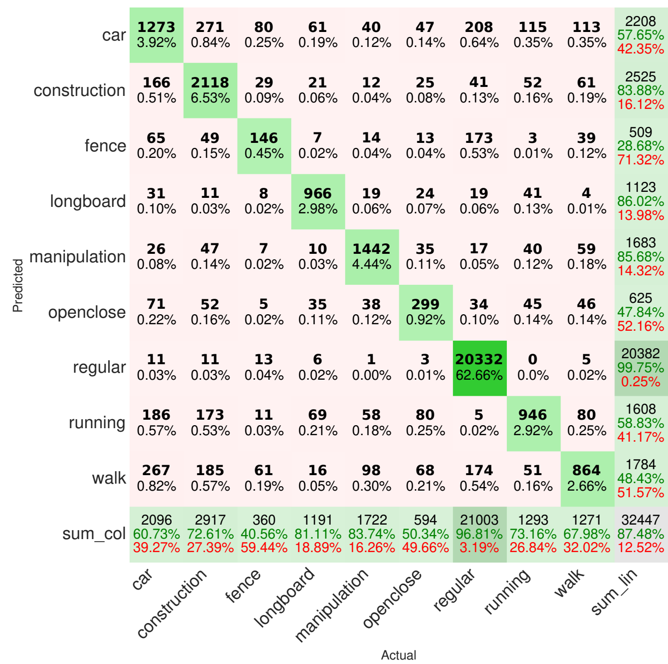

Fig. 4. Confusion matrix of the test set using the MLP. The MLP achieves substantially higher recall on *car* (∼70%) compared to the CNN (∼10%), though *fence* remains challenging with only ∼31% recall. The confusion between *car* and *construction* is reduced but still present.

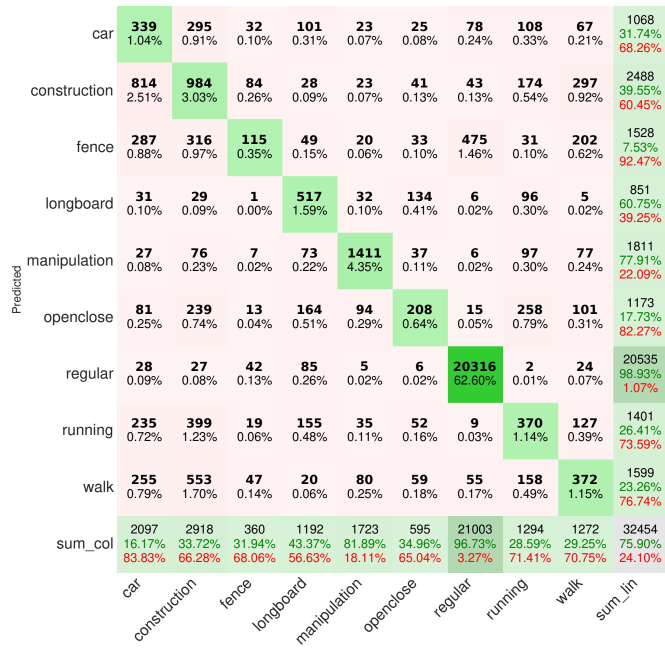

Fig. 5. Confusion matrix of the test set using the proposed MS-SNN in the 9-way (all 9 classes) 20-shot (20 samples from each) configuration. Despite using only 20 support samples per class, the model maintains >96% recall on *regular* while achieving more uniform performance across minority classes compared to the CNN.

Under this split, the bar for practical deployment is the 97%-level anomaly-detection performance already demonstrated, not the single-model 9-class score that would otherwise need to approach supervised-baseline accuracy. Closing that gap within one model is left to future work.

Because only 5–10 support samples per class are needed, labeling effort drops sharply compared to supervised retraining. New anomaly types are incorporated by simply adding support samples, addressing the *zero-day* anomaly case without weight updates. Combined with class-balanced episodic training, this also avoids the accuracy-inflation effect observed in the CNN/MLP results.

V. Conclusion

We presented an MS-SNN for fiber-sensing event recognition that combines five complementary similarity metrics with class-balanced episodic training. On a public DAS dataset, it reaches 97% accuracy and 98% recall for binary anomaly detection with only 5–10 support samples per class, matching specialized supervised approaches without extensive labeled training data, and handles novel event types without retraining. Our analysis also shows that standard accuracy can overestimate model utility by 18–24 percentage points on imbalanced sensing data, with balanced accuracy providing a more honest view.

The main limitation is that the study is confined to a single Φ-OTDR deployment, so the cross-system generalization argued for in Sec. II is motivated rather than empirically validated in this paper. Fine-grained 9-class classification with a single model also remains limited and is better addressed by a two-stage deployment. Future work includes cross-modality validation between DAS, SOP, and OTDR systems, threshold-sensitivity analysis via ROC/PR curves, integration with temporal models, and explainability techniques to support trustworthy AI in network automation.

References

[1] Y. Wang, H. Yuan, X. Liu, Q. Bai, H. Zhang, Y. Gao, and B. Jin, "A comprehensive study of optical fiber acoustic sensing," *IEEE Access*, vol. 7, pp. 85 821–85 837, 2019. DOI: 10.1109/ACCESS.2019.2924736.

[2] A. Tomasov, P. Zaviska, P. Dejdar, O. Klicnik, T. Horvath, and P. Munster, "Comprehensive dataset for event classification using distributed acoustic sensing (DAS) systems," *Scientific Data*, vol. 12, no. 1, p. 793, May 2025. DOI: 10.1038/s41597-025-05088-4.

[3] L. Sadighi, S. Karlsson, C. Natalino, L. Wosinska, M. Ruffini, and M. Furdek, "Deep learning for detection of harmful events in real-world, noisy optical fiber deployments," *Journal of Lightwave Technology*, vol. 43, no. 13, pp. 6092–6101, Jul. 2025. DOI: 10.1109/JLT.2025.3557748.

[4] G. A. Wellbrock, T. J. Xia, M.-F. Huang, S. Han, Y. Chen, T. Wang, and Y. Aono, "Explore benefits of distributed fiber optic sensing for optical network service providers," *Journal of Lightwave Technology*, vol. 41, no. 12, pp. 3758–3766, 2023. DOI: 10.1109/JLT.2023.3263795.

[5] G. Malik, I. C. Dipto, M. U. Masood, A. Ali, M. C. Mohamed, S. Straullu, S. K. Bhyri, G. M. Galimberti, J. Pedro, A. Napoli, W. Wakim, and V. Curri, "AI-driven fault prediction and restoration leveraging real-time SOP monitoring," in *2025 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, Sep. 2025, pp. 1–6. DOI: 10.23919/SoftCOM66362.2025.11197358.

[6] A. Tomasov, P. Zaviska, V. Spurny, P. Dejdar, P. Munster, T. Horvath, and O. Klicnik, "Enhancing perimeter protection using Φ-OTDR and CNN for event classification," in *28th International Conference on Optical Fiber Sensors (2023), Paper W4.39*, Nov. 2023, p. W4.39. DOI: 10.1364/OFS.2023.W4.39.

[7] L. Sadighi, C. Natalino, S. Karlsson, M. Ruffini, E. Kenny, L. Wosinska, and M. Furdek, "Generalizability of ML-based classification of state of polarization signatures across different bands and links," in *51th European Conference on Optical Communication (ECOC)*. IEEE, Oct. 2025, p. Th.02.01.2. DOI: 10.1109/ECOC66593.2025.11263096.

[8] S. Cruzes, "Failure management overview in optical networks," *IEEE Access*, vol. 12, pp. 169 170–169 193, 2024. DOI: 10.1109/ACCESS.2024.3498704.

[9] F. Musumeci and M. Tornatore, "Failure management in optical networks with ML: A tutorial on applications, challenges, and pitfalls [invited]," *Journal of Optical Communications and Networking*, vol. 17, no. 8, pp. C144–C155, Aug. 2025. DOI: 10.1364/JOCN.551910.

[10] M. Bublin, "Event detection for distributed acoustic sensing: Combining knowledge-based, classical machine learning, and deep learning approaches," *Sensors*, vol. 21, no. 22, p. 7527, Jan. 2021. DOI: 10.3390/s21227527.

[11] R. Gu, Z. Yang, and Y. Ji, "Machine learning for intelligent optical networks: A comprehensive survey," *Journal of Network and Computer Applications*, vol. 157, p. 102576, May 2020. DOI: 10.1016/j.jnca.2020.102576.

[12] G. Koch, R. Zemel, and R. Salakhutdinov, "Siamese neural networks for one-shot image recognition," in *ICML Deep Learning Workshop*, vol. 2, 2015.

[13] C. Natalino, A. Udalcovs, L. Wosinska, O. Ozolins, and M. Furdek, "One-shot learning for modulation format identification in evolving optical networks," in *OSA Advanced Photonics Congress (AP) 2019*, 2019, p. JW4A.2.

[14] Y. Gao, B. Guo, Y. Zhou, Y. Ma, K. Yan, and S. Huang, "Fault tracing based on siamese neural network for optical networks," in *2024 IEEE Opto-Electronics and Communications Conference (OECC)*, Jun. 2024, pp. 1–3. DOI: 10.1109/OECC54135.2024.10975324.

[15] C. Natalino, F. Monteiro Pessoa, and P. Monti, "A unified Siamese learning framework for zero-day anomaly detection and classification in optical networks," in *Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2026, p. M3A.2.

[16] K. Kikuchi, "Fundamentals of coherent optical fiber communications," *Journal of Lightwave Technology*, vol. 34, no. 1, pp. 157–179, Jan. 2016. DOI: 10.1109/JLT.2015.2463719.

[17] C. Pendão and I. Silva, "Optical fiber sensors and sensing networks: Overview of the main principles and applications," *Sensors*, vol. 22, no. 19, Oct. 2022. DOI: 10.3390/s22197554.

[18] Y. Shi, Y. Wang, L. Zhao, and Z. Fan, "An event recognition method for Φ-OTDR sensing system based on deep learning," *Sensors*, vol. 19, no. 15, Aug. 2019. DOI: 10.3390/s19153421.

[19] L. E. Kruse, S. Kuhl, A. Dochhan, and S. Pachnicke, "Experimental investigation of machine-learning-based soft-failure management using the optical spectrum," *Journal of Optical Communications and Networking*, vol. 16, no. 2, pp. 94–103, Feb. 2024. DOI: 10.1364/JOCN.500930.

[20] K. Murakami, C. Kito, and K. Toge, "Anomaly detection for telecom aerial optical cables with variational auto encoder leveraging DAS data," *Journal of Lightwave Technology*, pp. 1–9, 2026. DOI: 10.1109/JLT.2026.3652537.

[21] A. Knapinska and M. Furdek, "Experimental analysis of adaptive ML classifiers for dynamic detection of emerging physical-layer attacks," in *2025 European Conference on Optical Communications (ECOC)*, Sep. 2025, p. W.02.01.180. DOI: 10.1109/ECOC66593.2025.11263047.

[22] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. F. Wang, and J.-B. Huang, "A closer look at few-shot classification," in *International Conference on Learning Representations*, Sep. 2018.

[23] X. Luo, H. Wu, J. Zhang, L. Gao, J. Xu, and J. Song, "A closer look at few-shot classification again," in *Proceedings of the 40th International Conference on Machine Learning*. PMLR, Jul. 2023, pp. 23 103–23 123. [Online]. Available: https://proceedings.mlr.press/v202/luo23e.html

[24] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[25] H. V. Luong, N. Deligiannis, R. Wilhelm, and B. Drapp, "Few-shot classification with meta-learning for urban infrastructure monitoring using distributed acoustic sensing," *Sensors*, vol. 24, no. 1, p. 49, Jan. 2024. DOI: 10.3390/s24010049.

[26] Y. Wu, C. Zhao, J. Jiang, Y. Zhao, J. Wang, and H. Wu, "Few-sample learning for fiber-optic distributed acoustic sensors (DAS) in cross scenes," in *First Conference on Distributed Optical Fiber Sensing Technology and Applications (DOFS)*, vol. 13654, May 2025, pp. 205–210. DOI: 10.1117/12.3059262.