



A Unified Siamese Learning Framework for Zero-Day Anomaly Detection and Classification in Optical Networks

Downloaded from: <https://research.chalmers.se>, 2026-04-14 12:08 UTC

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

Natalino Da Silva, C., Pessoa Monteiro, F., Monti, P. (2025). A Unified Siamese Learning Framework for Zero-Day Anomaly Detection and Classification in Optical Networks. Proceedings of the Optical Fiber Communication Conference (OFC) 2026

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

A Unified Siamese Learning Framework for Zero-Day Anomaly Detection and Classification in Optical Networks

Carlos Natalino ^{1,*}, Flávia P. Monteiro ² and Paolo Monti ¹

¹Department of Electrical Engineering, Chalmers University of Technology, 412 96 Gothenburg, Sweden

²Federal University of Western Pará (UFOPA), 68040-255 Santarém, Pará, Brazil

*carlos.natalino@chalmers.se

Abstract: A multi-similarity Siamese neural network unifies zero-day anomaly detection and one-shot classification in optical networks, achieving over 99% accuracy and instant adaptability across lightpaths and unseen anomaly types without any retraining.

1. Introduction

Optical networks are increasingly complex systems, where anomalies arise from the interplay of heterogeneous devices, configurations, physical impairments, or external factors such as infrastructure degradation [1]. Reliable operation therefore requires flexible mechanisms for anomaly detection and classification. However, the massive volume of monitoring data makes manual oversight and threshold-based methods impractical for network-wide deployments, driving operators toward artificial intelligence & machine learning (AI/ML)-based automation [1,2]. Yet, models trained for specific lightpaths or device/network configurations often fail to generalize, calling for invariant, network-wide representations capturing the underlying behavior of the optical infrastructure.

A second challenge is the integration of anomaly detection and classification. Traditional approaches treat detection (i.e., identifying whether an anomaly exists) and classification (i.e., identifying the anomaly type or class) as separate stages. For example, [3] proposed a long short-term memory (LSTM) autoencoder (AE) for anomaly detection followed by an LSTM classifier for soft-failure identification, while [4] combined an AE for detection with an attention-based bidirectional gated recurrent unit (GRU) for fault diagnosis. Other works explored domain adaptation and transfer learning to reduce the need for labeled samples [5]. Although effective, these methods rely on rigid decision boundaries and offer limited transferability, making them unsuitable for dynamic conditions.

An additional difficulty is the zero-day anomaly problem, caused by previously unseen anomaly types not part of the training data. The anomaly landscape evolves continuously with new technologies, traffic dynamics, or environmental conditions, making static models quickly obsolete. While anomalies are rare, this scarcity of labeled data hinders robust model development. Supervised classifiers struggle with novel anomalies without complete retraining, and unsupervised methods can detect but not classify them. This motivates few-shot or one-shot learning models capable of recognizing and classifying a new anomaly from its first occurrence (i.e., day zero) without retraining. Recent studies have explored the use of Siamese neural networks (SNNs) and similarity-based learning for such low-sample regimes. For example, [6] combined an SNN with pattern mining to handle imbalanced anomaly data and perform one-shot clustering of new anomaly types, while [7] applied SNNs to modulation format classification, achieving strong generalization to unseen classes, but using only a single similarity metric. However, a single metric may be insufficient to capture the multidimensional dependencies among optical performance parameters, limiting generalization across diverse network conditions.

This paper introduces a multi-similarity Siamese neural network (MS-SNN) framework for zero-day anomaly detection and one-shot classification in optical networks. The framework: (i) generalizes across different lightpaths and network conditions through multi-similarity feature learning, (ii) integrates detection and classification into a unified learning process, and (iii) enables one-shot classification of unseen anomalies from their first observation. In contrast to SNN that rely on a single similarity metric, the proposed MS-SNN framework captures a richer representation of relationships among optical performance parameters, thus improving model robustness to unseen anomaly types. Experiments on an open dataset [3] show over 99% accuracy in one-shot classification scenarios, demonstrating a major step toward adaptive, network-wide anomaly intelligence in future optical infrastructures.

2. Zero-day anomaly detection and classification based on MS-SNNs

Our proposed framework employs a MS-SNN designed for generalizable zero-day anomaly detection in optical networks. Instead of learning a direct classification boundary, the SNN learns a similarity function over pairs of




Fig. 1. Architecture of the proposed multi-similarity Siamese neural network (MS-SNN) with multi-similarity metrics: two samples share the same embedding/encoder segment.

samples, i.e., a reference from a known class and a sample of interest. The SNN comprises three main segments: an embedding/encoder, a multi-similarity head, and a similarity computation module. Each sample is processed by the shared embedding/encoder, which extracts a high-dimensional feature representation. These representations are compared through a multi-similarity head, which computes complementary distance metrics, followed by a fully-connected layer that outputs a similarity score indicating whether the two samples belong to the same class.

Fig. 1 illustrates the proposed framework during operation. An incoming optical performance monitoring (OPM) sample is compared against a gallery of reference samples, each representing a known anomaly class. If the similarity score between the new sample and all known reference samples falls below a calibrated threshold, the system flags it as a new, unknown (or zero-day) anomaly class, which can be labeled afterwards by specialists. The first sample of this new class is then added to the reference gallery, allowing subsequent samples to be classified without retraining. This mechanism enables one-shot classification and allows the model to continuously expand its knowledge base. Although classification complexity scales linearly with the number of classes, parallel inference on modern artificial intelligence (AI) accelerators completes all comparisons within milliseconds, making the approach practically feasible for real-time operations.

The embedding/encoder consists of fully-connected layers mapping OPM parameters into a high-dimensional feature space. The multi-similarity head computes four similarity/distance measures (i.e., L1 (Manhattan), L2 (Euclidean), cosine similarity (CoS), and element-wise product) to capture different relationships between feature vectors. The similarity module then combines these measures through additional fully connected layers with learnable weights, enhancing the model's ability to capture non-linear correlations among OPM parameters.


The network is trained using contrastive learning with a binary cross-entropy loss function, reframing the multi-class problem into a binary decision: whether two samples are similar or not. During training, sample pairs from the same class are labeled as similar (1), while pairs from different classes are labeled as dissimilar (0). The objective of the training process is to minimize the binary cross-entropy loss, which effectively teaches the network to generate embeddings that are close together for samples of the same class and far apart for samples from different classes. This strategy is effective in data-scarce scenarios (i.e., anomaly management in optical networks) because it combinatorially expands the number of training examples from a small labeled set, allowing the model to learn a robust similarity metric without a large number of instances for each class. By learning a general notion of similarity rather than fixed decision boundaries, the multi-similarity Siamese neural network (MS-SNN) remains independent of the total number of classes and can generalize to unseen anomalies without modification.

Once trained, the same MS-SNN supports both anomaly detection and classification. For zero-day anomaly classification, only one reference sample per class is required. The class of a new sample is determined by averaging and/or thresholding its similarity scores across all reference samples, enabling efficient one-shot learning and strong generalization across network conditions.


3. Evaluation of the Proposed Framework

We use the dataset from [3] to validate the proposed multi-similarity Siamese neural network (MS-SNN) framework. It considers a 28-node, 41 links optical backbone network with inline amplifiers every 80 km and coherent receivers with OPM capability, providing quality of transmission (QoT) parameters. Each OPM sample includes five features: lightpath length, laser current, lightpath received power, optical signal-to-noise ratio (OSNR), and bit error rate (BER). Samples were collected from 756 lightpaths (one between every source-destination node pair) using the shortest path. The dataset was generated by simulating three types of soft failures dominant in backbone optical networks: external cavity laser (ECL) malfunction, failures related to increasing threshold current; Erbium-doped fiber amplifier (EDFA) malfunction, failures due to decreasing pump laser power/gain of inline amplifiers; and nonlinear interference (NLI) soft failures caused by randomly increasing lightpath transmit power. The QoT data comprise 900 time samples (approximately 150 hours) per lightpath collected under normal and faulty conditions. At any given time, a lightpath can experience at most one failure type.

The MS-SNN architecture comprises 5 fully-connected layers with batch normalization, ReLU activation, and 0.3 dropout rate, with 128, 256, 128, 64, and 32 neurons, respectively. The resulting embeddings are used to



(a) Accuracy



(b) Confusion matrix for MS-SNN

Fig. 2. Performance comparison between single-sample methods (MS-SNN, single-similarity SNNs, and 1-NN) and conventional ML models (naive Bayes (NB), gradient boosting (GB), random forest (RF), and linear/radial basis function (RBF) SVMs) trained on all four anomaly classes.

compute the four distance metrics discussed earlier: Manhattan, Euclidean, element-wise product, and cosine similarity. These are concatenated and fed to a fully-connected layer with softmax activation, followed by an output neuron with sigmoid activation. For training, we used 1,000 samples from all lightpaths and two classes: 0 (normal) and ECL malfunction (1). This deliberately limited training set allows evaluating the model's zero-day detection and classification capability. Testing used 10,000 samples from each of the four classes to simulate zero-day scenarios, comprising samples from all lightpaths. The network was trained for 100 epochs using the Adam optimizer (learning rate of 0.001, batch size of 64).

To assess the benefit of multi-similarity learning, we benchmark the proposed MS-SNN against four variations of the same SNN architecture, each using a single similarity metric. The MS-SNN was also compared with the nearest neighbor (NN) algorithm (1-NN), which also relies on one sample per class, and with traditional ML models (NB, GB, RF, and linear and RBF SVMs) that require labeled data for all classes. These baseline ML models were trained with 1,000, and tested over 10,000 samples per class. Unlike the proposed framework, they may not detect unseen anomalies, and cannot classify unseen anomaly types, being limited to closed-set classification.

Fig. 2 summarizes the results. The proposed MS-SNN achieves 99.12% accuracy, outperforming single-metric SNNs and 1-NN (97.6%) (Fig. 2(a)). Among the single-metric variants, the one with CoS performs best yet below 99%, followed by L1 with 96.6%. Surprisingly, the SNN element-wise product as similarity metric only achieves 25.9% accuracy, which is very close to the theoretical random classifier performance for four classes (i.e., 25%). This confirms that combining multiple similarity measures yields richer, more discriminative representations and stronger generalization to unseen anomaly types. Traditional ML methods (NB, GB, and RF) achieve slightly higher accuracy (around 99.8%), but rely on labeled data for all classes, defeating the purpose of zero-day classification. This emphasizes the strength of MS-SNN: high accuracy with minimal supervision and no retraining. The confusion matrix in Fig. 2(b) shows 97% anomaly detection accuracy, i.e., around 3% of false negatives, where the failures are misclassified as normal operating conditions. Notably, the MS-SNN perfectly identifies normal operating states, i.e., it produces no false positives, avoiding unnecessary alarm overhead.

4. Final remarks

This work proposed an MS-SNN-based framework that achieves zero-day anomaly detection and one-shot classification with over 99% accuracy and no false positives, paving the way for scalable, adaptive anomaly management in future optical networks. Future works may explore the suitability of the proposed approach over other anomaly detection and classification scenarios, including anomaly identification, and localization.

Acknowledgments: This work has been supported by the Horizon Europe ECO-eNET project with funding from the SNS JU under grant agreement No. 101139133.

References

1. F. Musumeci, *et al.*, *J. Opt. Commun. Netw.* **17**, C144–C155 (2025). DOI: [10.1364/JOCN.551910](https://doi.org/10.1364/JOCN.551910).
2. X. Chen, *et al.*, *IEEE Commun. Mag.* **60**, 88–94 (2022). DOI: [10.1109/MCOM.003.2200110](https://doi.org/10.1109/MCOM.003.2200110).
3. S. Ghosh, *et al.*, *Comput. Networks* **262**, 111159 (2025). DOI: [10.1016/j.comnet.2025.111159](https://doi.org/10.1016/j.comnet.2025.111159).
4. K. Abdelli, *et al.*, *J. Opt. Commun. Netw.* **14**, 365–375 (2022). DOI: [10.1364/JOCN.451289](https://doi.org/10.1364/JOCN.451289).
5. F. Musumeci, *et al.*, *J. Opt. Commun. Netw.* **14**, A91–A100 (2022). DOI: [10.1364/JOCN.438269](https://doi.org/10.1364/JOCN.438269).
6. Y. Gao, *et al.*, in *Proc. of OECC*, (2024). DOI: [10.1109/OECC54135.2024.10975324](https://doi.org/10.1109/OECC54135.2024.10975324).
7. C. Natalino, *et al.*, in *Proc. Optica APC*, (2019). P. JW4A.2. ISBN: 978-1-943580-64-4.