Online Discovery and Incremental Learning of Unknown Soft Failures in Optical Networks

Azarm Yeganehfallah[1], Emilio Paolini[1], Andrea Sgambelluri[1], Luca Valcarenghi[1]
[1] TeCIP Institute, Scuola Superiore Sant'Anna, Pisa, Italy
azarm.yeganehfallah@santannapisa.it

***Abstract*—We propose an online adaptive learning framework for the detection and discovery of previously unknown soft failures in optical networks. The method combines autoencoder-based feature extraction with a one-vs-rest multiclass classifier and confidence-driven novelty detection. New failure behaviors are autonomously identified and incrementally learned during live operation. Experimental validation on a laboratory optical network testbed demonstrates reliable new-class discovery and up to 97% classification accuracy.**

***Index Terms*—online learning, soft failure, optical network.**

I. Introduction

Optical network soft failures (SFs) emerge as gradual and unexpected performance degradations that may precede hard failures. In recent years, machine learning techniques have been widely applied to SF detection and diagnosis in optical networks, primarily through supervised classification approaches trained on predefined failure categories [1]. Although these methods achieve high accuracy when comprehensive labeled datasets are available, they inherently assume a closed set of known failures and therefore fail to generalize to previously unseen or evolving conditions. To relax this assumption, unsupervised and semi-supervised methods, often based on autoencoders (AEs) or clustering, have been proposed to detect anomalies without explicit failure labels [2]. However, such approaches typically treat all anomalies as a single class and do not support failure-category identification or adaptation of a structured model during operation. More recent adaptive and online learning solutions address pattern change in data by periodically updating the model [3], but still assume a fixed class set and do not explicitly target the discovery of unknown failure.

These limitations motivate adaptive monitoring mechanisms capable of autonomously detecting and learning new failures during live operation. In this work, we address this challenge by proposing a hybrid model that integrates AE with a one-vs-rest (OvR) multiclass classification strategy [4] implemented using a multilayer perceptron (MLP). The proposed model leverages reconstruction error and classification confidence for effective novelty detection, while enabling the incremental expansion of its knowledge base (KNB) to incorporate newly observed soft failure types. This design allows the system to continuously evolve and improve without requiring prior knowledge of all possible failure conditions. The main contributions of this work are summarized as follows: (i) an online learning model for the detection and classification of previously unseen SFs; (ii) a novelty detection mechanism driven by reconstruction error and classification confidence; (iii) an incremental adaptation strategy that enables continuous model evolution; (iv) experimental validation on a real optical network testbed.

This work is supported by the Chips Joint Undertaking (JU), European Union (EU) HORIZON-JU-IA, under grant agreement No. 101140087 (SMARTY).

II. Proposed Online Learning Model

As a big picture, the core conceptual part of our proposed model is depicted in Fig. 1. However in detail, the framework includes three phases: pre-initialization, initial training, and live monitoring. In the pre-initialization phase, a balanced and normalized training dataset comprising $\boldsymbol{n}$ known classes is generated, referred to as the original KNB. Based on this training set, an AE is then optimized by searching over a range of hyperparameters (i.e., learning rate, batch size, architecture, and regularization). The hyperparameters that yield the lowest validation loss are selected. This procedure is performed only once to align the AE with the dataset type.

In the initial phase, the same original KNB is used to train the AE to capture representative latent features while preserving essential characteristics and reducing dimensionality. In addition, the mean squared reconstruction error (MSRE), defined as MSRE $= \frac{1}{n}\sum_{i=1}^{n} |\mathbf{x}i - \hat{\mathbf{x}}i|^2$, is computed on the training set. An updatable threshold, denoted as $\boldsymbol{Trsh}_{\mathbf{rec}}$, is then determined at a selected percentile of the MSRE distribution, i.e., $Trsh_{\text{rec}}$ = ($Percentile_{\text{rec}}$ * MSRE). Subsequently, an OvR multiclass classification strategy is employed, where an independent MLP classifier is trained per class to discriminate it from all other classes. The class with the highest predicted probability is selected as the final decision for each sample. Based on these probabilistic outcomes, a probability threshold, denoted as $\boldsymbol{Trsh}_{\mathbf{prob}}$, is defined as the minimum probability observed across original KNB (excluding 5% extreme low to eliminate noise). This threshold represents the minimum confidence level required for the model to assign a class label to a sample and will be constant throughout the third phase. While $\boldsymbol{Trsh}_{\mathbf{rec}}$ is set at a high percentile to minimize false novelty detections, the probability threshold $\boldsymbol{Trsh}_{\mathbf{prob}}$ enforces conservative decision-making for class assignment.

Once the model is trained on original KNB, it begins processing live monitoring samples (i.e., third phase). To emulate live network monitoring, the entire dataset is processed sequentially, sample by sample. For each incoming sample, the

978-3-903176-78-2

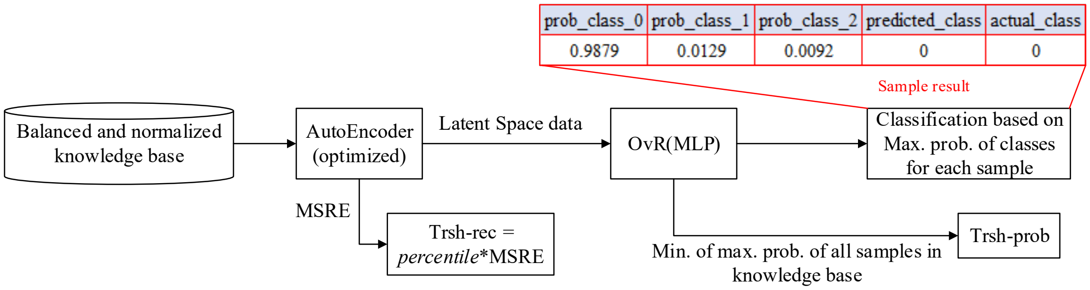

Fig. 1. The main pipeline of the online novelty detection and incremental learning model

classifier assigns the most probable class label. For samples originating from known classes, the model typically yields high-confidence predictions. In contrast, when an unknown condition e.g., new SFs occurs, the model's confidence in predicting a class falls below $\boldsymbol{Trsh}_{\mathbf{prob}}$. Such snapshots are cached and further examined using $\boldsymbol{Trsh}_{\mathbf{rec}}$ to distinguish genuinely novel events from cases where the model's prediction is correct but of low confidence.

New class of SF diagnosis is confirmed when both of the following conditions are satisfied: (1) more than $\boldsymbol{k}$ consecutive snapshots have probabilities below $\boldsymbol{Trsh}_{\mathbf{prob}}$, and (2) at least $\boldsymbol{m}$ number of these snapshots exhibit reconstruction errors exceeding $\boldsymbol{Trsh}_{\mathbf{rec}}$. In this case, a new class of SF is declared, prompting the retraining of the model to incorporate the new class. Conversely, if only the first condition is satisfied, the model is fine-tuned to improve prediction confidence and preventing the model to downturn respect to temporal pattern change. Hence, the retraining and fine-tuning phases are largely similar, differing only in labeling: in both cases, the $\boldsymbol{k}$ cached snapshots are incorporated into the KNB, with retraining using a new class label and fine-tuning using the predicted class label. Then, to ensure balanced representation across the KNB, the $\boldsymbol{k}$ cached samples are augmented to match the size of the training set by adding Gaussian noise. After that, following the same procedure as in the initial setup, MSRE and consequently $\boldsymbol{Trsh}_{\mathbf{rec}}$ is updated, and the tuned parts replace the initial one in the model. The $\boldsymbol{k}$ consecutive samples are important to distinguish noise from persistent outliers indicative of emerging SF class. The supporting parameter $\boldsymbol{m}$ is set to $\boldsymbol{k}$-2, to have integrated sensitivity analysis, while requiring that most of the $\boldsymbol{k}$ low-confidence samples also exhibit high reconstruction errors. In timing point of view, each snapshot corresponds to a monitoring interval of $\boldsymbol{s}$ seconds and spans $\boldsymbol{k}$ consecutive snapshots. Therefore, a total duration of $\boldsymbol{s}\times\boldsymbol{k}$ seconds is required to confidently identify the onset of a new SF class, ensuring it is not mistaken for random noise. This aligns with the typical behavior of SFs in optical networks, which manifest as gradual degradations, rather than exhibiting transient pattern.

III. Experimental Evaluation

The proposed framework is evaluated using telemetry data collected through a Kafka-based monitoring system [5] deployed on a laboratory optical network testbed, as illustrated in Fig. 2(b). The setup comprises two transmit (TX) and two receive (RX) coherent muxponders at $100\mathrm{Gb/s}$, operating on adjacent DWDM channels at $193.1, \mathrm{THz}$ (Ch1) and $193.2, \mathrm{THz}$ (Ch2). The transceivers are interconnected through a four-span optical link, with each span extending over 70 km and equipped with an EDFA to compensate for power loss. To enable controlled fault injection, a Failure Actuator (FA) incorporating a Wavelength Selective Switch (WSS) is deployed within each span. Each FA is capable of inducing three distinct SF modes corresponding to degradations affecting Ch1, Ch2, or both channels simultaneously.

A. Evaluation Criteria

Performance is evaluated along three complementary dimensions: (i) the correct identification of previously unseen failure classes, (ii) the correspondence between the discovered

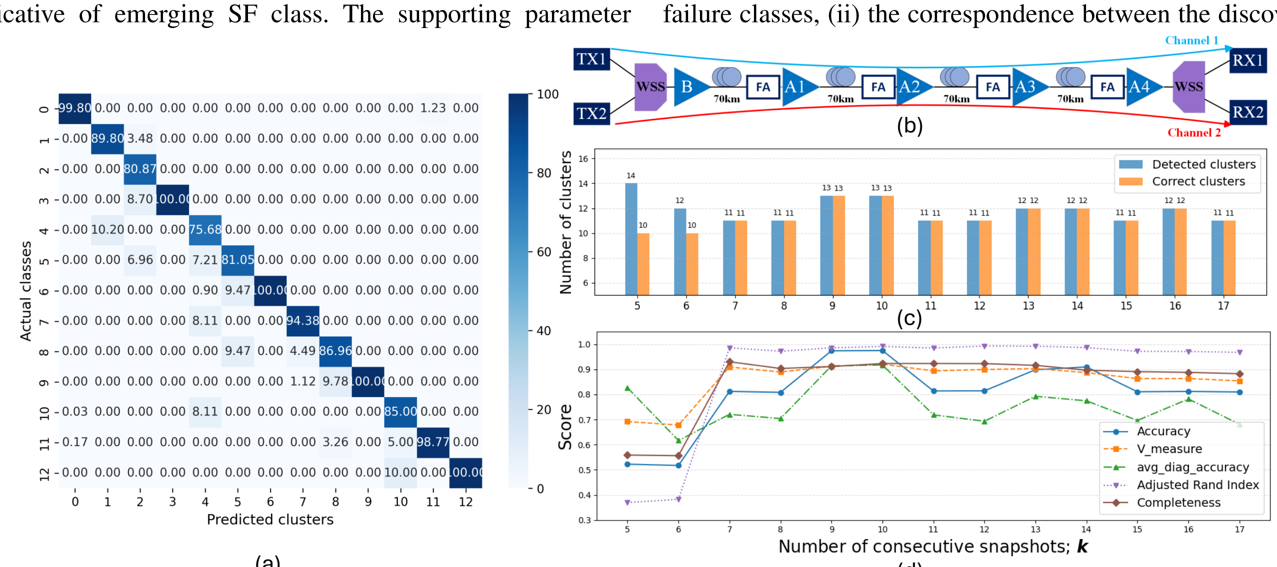

Fig. 2. (a) Confusion matrix with $\boldsymbol{k}$ = 10, (b) Considered testbed for data collection, (c) Correct identification and correspondence between the discovered clusters and the true failure labels as a function of $\boldsymbol{k}$, (d) Evaluation scores as a function of $\boldsymbol{k}$.

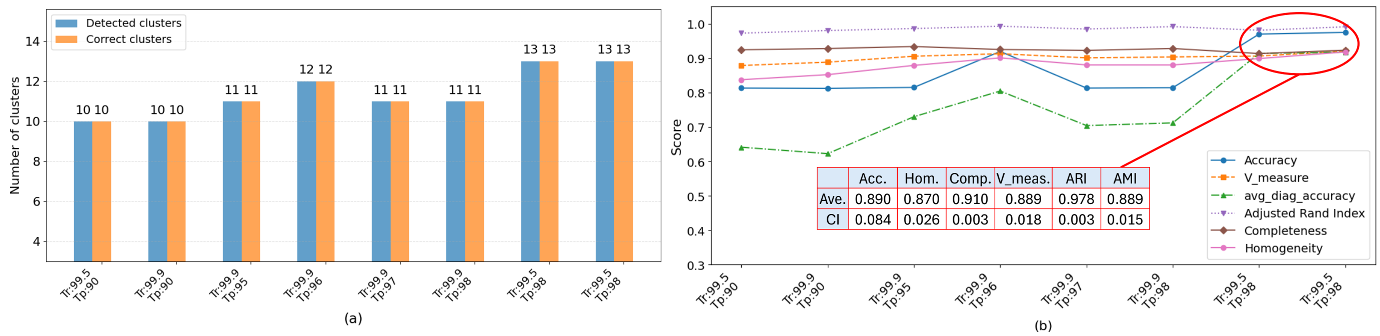

Fig. 3. (a) Correct identification and correspondence between the discovered clusters and the true failure labels (b) Evaluation scores both as a function of $\boldsymbol{Trsh}_{\mathbf{rec}}$ percentile (Tr) and $\boldsymbol{Trsh}_{\mathbf{prob}}$ (Tp) plus 95% confidence interval of the metrics at the best values of the thresholds are provided in the table

clusters and the true failure labels, and (iii) the classification accuracy after incremental model adaptation. Two types of metrics are used, for classification and clustering. To evaluate classification, accuracy (i.e., the ratio of *Correct predictions* to *Total predictions*) and Average Diagonal Accuracy (ADA) (i.e., the mean accuracy of correctly matched clusters in the confusion matrix) are elaborated. Clustering evaluation metrics that are implemented, measure the quality of grouping unlabeled data by comparing predicted clusters to ground truth. Key metrics include Homogeneity (i.e., each cluster contains only members of a single class), Completeness (i.e., all members of a class assigned to the same cluster), V-measure (i.e., harmonic mean of the two), Adjusted Rand Index (ARI) (i.e., pairwise agreement adjusted for chance), and Adjusted Mutual Information (AMI) (i.e., agreement based on information theory).

B. Experimental Procedure and Parameters Setting

Telemetry measurements are sampled every four seconds (i.e., $\boldsymbol{s} = 4$) over a monitoring period exceeding five hours, resulting in twelve distinct SF states in addition to one normal operating condition. These thirteen operational states are manually annotated and used exclusively for evaluation purposes. Each monitoring sample is represented by a 12-dimensional feature vector composed of BER and OSNR of Ch1 and Ch2, along with the input and output optical power levels of the four EDFAs along the link. This study relies on the same preprocessed dataset introduced in [6], ensuring consistency and reproducibility. During the initial training phase, only a subset of three classes out of the thirteen available (i.e., normal operating state, SF on span 1, CH1 and CH2) is used to train the model, with 30 samples per class. As shown in Fig. 2(c) and confirmed by the convergence of multiple evaluation metrics in Fig. 2(d), values of $\boldsymbol{k} = 9$–10 provide the most reliable performance. While at this value of $\boldsymbol{k}$ the best overall performance is obtained for $\boldsymbol{Trsh}_{\mathbf{prob}} = 98\%$ and $\boldsymbol{Percentile}_{\mathbf{rec}} = 99.5\%$, as shown in Fig. 3.

C. Results and Discussion

Both the novelty discovery capability of the proposed model and its long-term stability under continuous operation are assessed with introduced criteria and metrics. Fig. 2(a) reports the accuracy-based confusion matrix obtained for $\boldsymbol{k} = 10$ based on evaluation criteria (iii). An overall detection and clustering accuracy of 97% is achieved, with class-wise accuracies ranging between 75% and 100%, indicating robust discrimination across diverse failure conditions. Fig. 2(c) and (d) depict how the capability of the model in clustering changes as a function of $\boldsymbol{k}$ based on evaluation criteria (ii, iii). The number of detected clusters should coincide with the number of correctly identified ones and ideally with the true number of classes (i.e., 13). Deviations from this condition indicate either over-fragmentation or missed detections. Smaller values of $\boldsymbol{k}$ lead to premature triggering and class fragmentation, whereas larger values improve stability at the cost of delayed detection. A sensitivity analysis on both thresholds is presented in Fig. 3, confirming that conservative thresholding effectively balances detection sensitivity and robustness. Fig. 3 (a) evaluates the clustering capability of the model under different values of $\boldsymbol{Trsh}_{\mathbf{prob}}$ and $\boldsymbol{Percentile}_{\mathbf{rec}}$, proving that in the selected values, 13 correct classes are successfully detected and being clustered (meeting criteria ii and iii). In Fig. 3 (b) metrics such as V-measure and ARI confirm a strong alignment between the discovered clusters and the true failure labels. Furthermore, the ADA, highlights a consistent one-to-one correspondence between predicted clusters and actual classes. Finally, a table of average and 95% confidence intervals per each metric is added to Fig. 3 (b).

IV. Conclusion

This work presents an online adaptive learning framework for the detection and discovery of unknown SFs in optical networks. By combining AE-based representation learning with confidence-driven multiclass classification, the proposed approach autonomously identifies emerging failure behaviors and incrementally adapts its KNB during live operation. Experimental results on a real optical network testbed demonstrate robust new-class discovery and sustained classification accuracy of up to 97%.

References

[1] Xiaoliang Chen et al. "Automating Optical Network Fault Management with Machine Learning". In: *IEEE Communications Magazine* (2022).

[2] Xiaoliang Chen et al. "Self-Taught Anomaly Detection With Hybrid Unsupervised/Supervised Machine Learning in Optical Networks". In: *J. Lightwave Technol.* (2019).

[3] Sergio Cruzes. "Failure Management Overview in Optical Networks". In: *IEEE Access* (2024).

[4] Ryan Rifkin and Aldebaro Klautau. "In Defense of One-Vs-All Classification". In: *Journal of Machine Learning Research* (2004).

[5] Andrea Sgambelluri et al. "Reliable and scalable Kafka-based framework for optical network telemetry". In: *J. Opt. Commun. Netw.* (2021).

[6] Azarm Yeganehfallah et al. "Effectiveness of Confidentiality-Preserving Clustering Algorithms for Soft Failure Detection in Optical Networks". In: *IEEE HPSR conference*. Pisa, Italy, 2024, pp. 87–92.