Spectrum Anomaly Detection for Optical Network Monitoring using Deep Unsupervised Learning

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

(article starts on next page)

Spectrum Anomaly Detection for Optical Network Monitoring using Deep Unsupervised Learning

Carlos Natalino, *Member, IEEE,* Aleksejs Udalcovs, *Member, IEEE,* Lena Wosinska, *Senior Member, IEEE,* Oskars Ozolins, and Marija Furdek *Senior Member, IEEE, OSA*

Abstract—Accurate and efficient anomaly detection is a key enabler for the cognitive management of optical networks, but traditional anomaly detection algorithms are computationally complex and do not scale well with the amount of monitoring data. Therefore, we propose an optical spectrum anomaly detection scheme that exploits computer vision and deep unsupervised learning to perform optical network monitoring relying only on constellation diagrams of received signals. The proposed scheme achieves 100% detection accuracy even without prior knowledge of the anomalies. Furthermore, operation with encoded images of constellation diagrams reduces the runtime by up to 200 times.

Index Terms—Deep unsupervised learning, anomaly detection, constellation diagram, autoencoder, optical network monitoring

I. INTRODUCTION

The proliferating high-performance network services and applications require increasingly dynamic, flexible and autonomous optical networks. Optical network telemetry [2] and machine learning (ML)-based approaches [3] are considered critical enablers of quick and reliable network (re)configuration through state sensing, analyses and learning from ever-larger optical performance monitoring (OPM) datasets. Among the multitude of possible ML applications in autonomous optical network operation, anomaly detection has particular importance thanks to its ability to spot an anomaly without specific knowledge of its signature. This can be used to trigger a response without depending on human expert interpretation and action [4], [5].

However, optical equipment collects diverse numerical/categorical data, furnishing a fragmented set of ML solutions that need different inputs for each objective. In this context, constellation diagrams show potential to become an interoperable input for a wide range of models assessing optical channel quality [6]–[9]. Constellation diagrams can be interpreted by both humans and ML algorithms, which facilitates the checking and validation of ML decisions by specialists. Nonetheless, anomaly detection using constellation

diagrams is challenging due to their high dimensionality. Pertinent algorithms usually require traversing the entire dataset, which builds up complexity [4], [10]–[12].

While [6]–[9] combine computer vision and supervised learning approaches to estimate the signal imperfections and quality of transmission (QoT), this letter demonstrates, for the first time, the use of computer vision and deep unsupervised learning (UL) for efficient and accurate anomaly detection. To reduce complexity, we develop an autoencoder that compresses constellation diagram images by extracting the feature information. In such a way, we decrease the dimensions of processed data before they are fed as input to the anomaly detection algorithm, yielding up to a 200-fold runtime reduction and improving the accuracy. With no prior knowledge, the proposed approach detects 100% of abnormalities imposed on the optical channel by analyzing only constellation diagrams. It outperforms one-class support vector machine (OCSVM), a conventional UL algorithm considered efficient in anomaly detection [12], [13], as well as density-based spatial clustering of applications with noise (DBSCAN).

II. DEEP UL FOR ANOMALY DETECTION

Figure 1 shows the architecture of the deep UL approach for anomaly detection proposed in this work. The approach combines the use of a convolutional neural network (CNN) and an anomaly detection algorithm to detect spectrum anomalies from the constellation diagrams.

A. Background

Constellation diagrams have been used for various OPM tasks, such as modulation format identification [6], [9] and optical signal-to-noise ratio (OSNR) estimation [7]. This indicates the potential of computer vision models in optical networks, allowing the use of a unified representation of an optical channel to identify/estimate different performance metrics. Moreover, such models can be combined to provide a broad assessment of the optical channel quality with a single model inference, i.e., provide multiple output values for the same input, potentially enhancing the efficiency of network management. However, all of the previous works that use computer vision algorithms to assess performance of optical networks [6]–[9] rely on supervised learning and thus require a dataset with previously observed samples clearly identifying the inputs and outputs of the problem. Due to the relatively infrequent occurrences of anomalies, the difficulties of collecting representative datasets, the potentially evolving
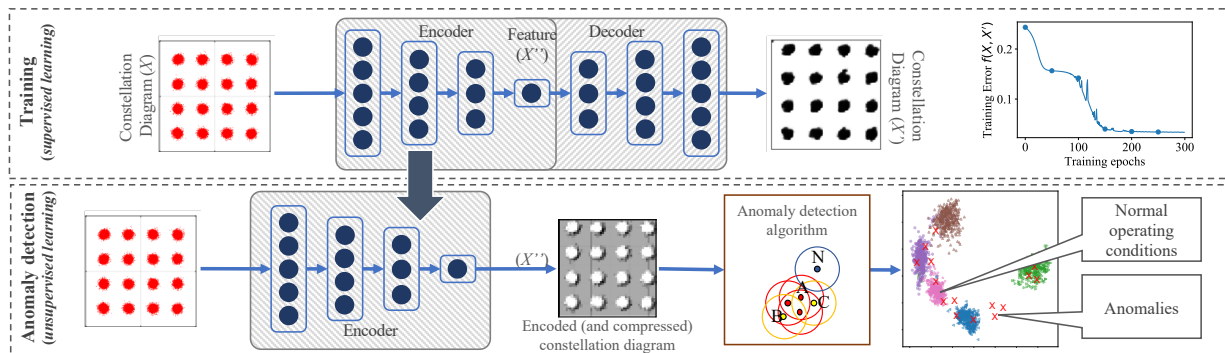
Fig. 1: The proposed deep UL approach for anomaly detection. Supervised learning is used to train an autoencoder (top). Anomaly detection is performed over encoded constellation diagrams, obtained by the encoder part (bottom).

anomaly landscape, and the time-consuming task of dataset labeling, it becomes impractical to rely solely on supervised learning algorithms for anomaly detection [4].

Unsupervised learning algorithms, on the contrary, learn from the data without prior knowledge, and have the potential to enable truly autonomous optical network monitoring and quality assurance. Algorithms such as the DBSCAN [10] have been successfully applied to detect anomalies [4] and jamming attacks [5] in optical networks. However, the complexity of DBSCAN, when considering Euclidean distance between the samples, is known to increase with the number of samples n and features (or dimensions) d in the order of $\Theta(n^2 \cdot d)$. Consequently, inputs to the algorithm must be critically sieved, especially when operating with images. Applying the density-based clustering directly on the obtained high-resolution images of constellation diagrams would lead to computationally-demanding and thus slow anomaly detection. To overcome this issue, we use a CNN-based autoencoder to encode (and decode) the features of a constellation diagram. In this way, the encoded representation contains only the most informative features, which allows the UL algorithm to concentrate on them, boosting the performance.

B. Autoencoder Training

The autoencoder used to reduce the dimensionality of constellation diagrams is a symmetrical neural network that consists of two parts—encoder and decoder. These are connected by a central layer responsible for the extraction and encoding the representation of the input. We use a supervised learning approach to train the autoencoder, as illustrated in the top part of Fig. 1. Since in our work constellation diagrams are represented as images, we use a CNN structure composed of convolutional and pooling layers. During the training, the original constellation diagram X is given as input, and it is also used to supervise the update process of the trainable CNN parameters in a way that the output diagram X' approximates the original X as closely as possible. After the training is complete, the encoder can be used to extract the encoded representation of any constellation diagram X'', concentrating the most informative features.

C. In-Operation Use of the Approach

The in-operation use of the proposed approach is illustrated in the bottom part of Fig. 1. Constellation diagrams are sampled at the optical receiver and preprocessed on-the-fly to make the images more suitable for ML algorithms by removing axes labels and grid lines. A preprocessed constellation diagram is input to the CNN-based encoder, and the encoded diagram is added to a dataset containing historical encoded constellation diagrams for the channel being analyzed. This dataset is input to the DBSCAN algorithm [10] that detects anomalies assuming that anomalous constellation diagrams are much rarer than the normal ones. Fig. 1 illustrates how DBSCAN works, i.e., by finding core samples (represented by the sample A) which have a minimum number of neighbors within a given radius, and border samples (samples B and C). When a sample does not have enough neighbors within the radius, it is considered an anomaly (sample N). The detected anomalies can be used to raise alarms, trigger corrective procedures, and/or be analyzed together with detailed data of the anomalous channel for (deeper) root cause analysis.

III. CONSIDERED SPECTRUM ANOMALIES

We consider anomalies caused by the interference between a neighboring channel and the channel under test (CUT, λ_1) in a wavelength division multiplexed (WDM) system that uses a 16-ary quadrature amplitude modulation (16QAM) for transmission at 32 Gbaud. The considered anomalies can appear due to, e.g., laser detuning due to ageing, temperature change, optical filter misalignment in optical nodes, or critical degradation, spectrum slot allocation problem, or even a deliberate attack inserted on an adjacent wavelength [14].

To assess the impact of the interferer on signal quality, we use a system configuration depicted in Fig. 2(a). We assume that *(i)* all channels ($\lambda_1 - \lambda_3$, and the interferer) use the same signal format; *(ii)* initially, all channels are equally spaced in the 50 GHz grid and have the same optical power level; *(iii)* the frequency interval between the interferer and the CUT is defined by parameter $\Delta f = \{50, 43.75, 37.5, 31.25\}$ GHz; and *(iv)* the power level difference between the interferer and the CUT is defined by parameter $\Delta p = \{-3, 0, 3, 6\}$ dB. We use the VPIphotonics Design Suite™ [15] to characterize the system performance. To consider different noise conditions for the
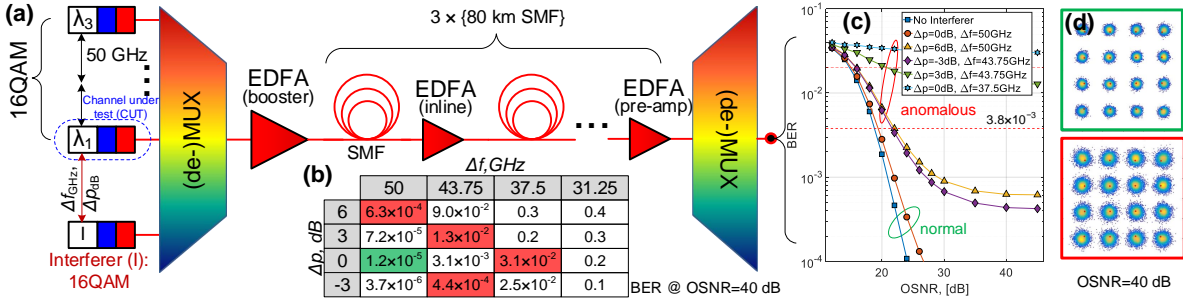
Fig. 2: (a) Simulation setup of the three-channel WDM system with the interferer; (b) BER detected for the CUT under assumed normal (green) and anomalous (red) conditions; (c) BER vs. OSNR curves for the selected cases; and (d) example of constellation diagrams: (top) no interferer and (bottom) impaired by the interferer for (Δp=-3 dB, Δf=43.75 GHz).

CUT, we set a specific OSNR value in the transmitter prior to inserting the interferer. To obtain statistically reliable results, each constellation diagram sample is generated by mapping a $2^{15} - 1$ long uniquely-seeded pseudorandom binary sequence (PRBS) to 16QAM symbols.

Figure 2 (b) shows the BER values of the CUT for 40 dB OSNR after it has been affected by the interferer characterized by different Δp and Δf values. The constellation diagrams for the selected normal and anomalous conditions are used to evaluate our approach. The normal operating conditions include two cases: when the interferer is not present; and when it is added to the system with appropriate settings, i.e., (Δp=0 dB, Δf=50 GHz), and, thus, behaves and is treated as a legitimate channel. We selected four cases among the remaining ones and labeled them as the anomalous operating conditions, considering the BERs, Δp and Δf values as well as other aspects. Finally, Fig. 2(c) characterizes the system performance with and without the presence of the interferer for an OSNR range of interest. Together with the provided examples of constellation diagrams (Fig. 2(d)), these illustrate the severity of the interferer-induced distortions.

IV. RESULTS AND DISCUSSION

The performance of the proposed autoencoder-assisted DBSCAN approach (DBSCAN-AE) is evaluated using data collected via simulations. 200 samples of constellation diagrams were collected under the two normal operating conditions, i.e., no interferer and (Δp=0 dB, Δf=50 GHz). These constellation diagrams are split 100:100 and used for training and validation when trying several combinations of CNN architectures and learning rates. The best architecture has 4 convolutional alternating with 4 pooling layers for both the encoder and the decoder, resulting in an encoded constellation diagram of 26×26 (676 pixels).

Two benchmarks are used to assess the benefits of the proposed approach: *(i)* OCSVM, a semi-supervised learning algorithm regarded as efficient in anomaly detection algorithm [12], [13] but requiring training to characterize normal operating conditions; and *(ii)* DBSCAN using un-encoded constellation diagrams (DBSCAN-U). OCSVM has three parameters: the kernel function (k), the up-
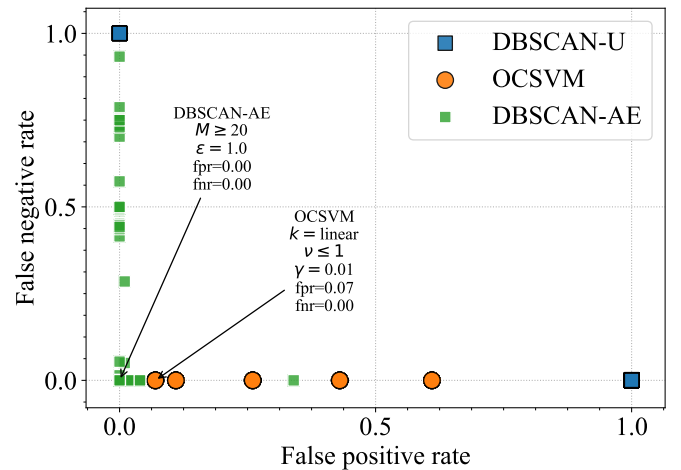

Fig. 3: Accuracy of the considered approaches and the settings giving the best performance.

per bound on the training error (ν) and γ which defines the coefficient of the kernel function (when applicable). We considered the following sets of the parameters: k={$linear, RBF, sigmoid$}, ν={0.01, 0.1, 0.3, 0.5, 0.7, 1} and γ={0.01, 0.1, 0.2, 0.5, 0.7, 1}. DBSCAN has two parameters: the minimum number of neighboring samples for a sample to be considered normal (M) and the neighborhood radius (ϵ). We considered the following sets of the parameters: M={1, 4, 8, 10, 12, 15, 20, 30, 40, 50} and ϵ={0.1, 0.5, 1, 1.5, 2, 2.5, 3, 4, 5, 10}. The algorithms are tested over all possible combinations of $<k, \nu, \gamma>$ values for OCSVM and $<M, \epsilon>$ values for DBSCAN. We perform 50 random experiments to obtain and average the results. During each experiment, we randomly select 100 constellation diagrams representing the normal operating conditions and 10 diagrams for each of the four anomalous conditions. By repeating this random selection procedure 50 times, we ensure that the obtained performance is not biased towards any particular part of the dataset. We use Python 3.7 and Scikit-learn version 0.21.2 [1] running on a workstation with a Core i9 9900X processor and 64 GB of RAM.

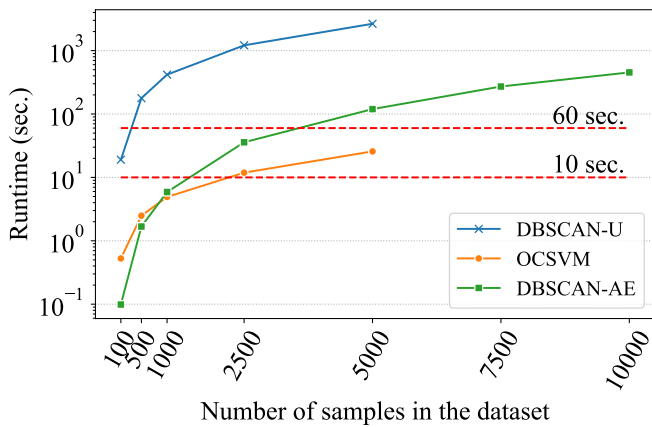To reveal which parameter settings yield the best accuracy

Fig. 4: Runtime of the considered approaches.

(i.e., lowest false positive and false negative rates), we performed a sensitivity analysis, considering different values of the parameters in the benchmarked algorithms. Fig. 3 shows the performance in terms of false positives and false negatives for different parameter settings, highlighting the best setting and performance achieved by each approach. DBSCAN-U is not able to detect the anomalies in any of the tested configurations, achieving either 100% false positives or 100% false negatives. This is explained by the sensitivity of the DBSCAN (and any other unsupervised learning algorithm) when operating with high-dimensional data. OCSVM performs better, achieving only 7% false positive rate with zero false negatives, but it requires a training procedure for each type of signal. On the contrary, our proposed DBSCAN-AE approach is able to achieve both zero false positive and false negative rates for configurations where $M \geq 20$ and $\epsilon = 1.0$.

Furthermore, we tested these anomaly detection approaches on two different signals: *(i)* 16QAM at 25 dB OSNR and *(ii)* 64QAM at 40 dB OSNR. Even without re-training the autoencoder, DBSCAN-AE achieves 1.9% false positive and 0% false negative rates for the 16QAM@25 dB signals. The corresponding numbers for the 64QAM@40 dB are 42% and 0%. In both cases, the DBSCAN-U fails at detecting anomalies. OCSVM, even when re-trained, performs slightly worse than the DBSCAN-AE for the 16QAM@25 dB, and similar for the 64QAM@40 dB signals. These results clearly identify DBSCAN-AE as an appropriate solution for detecting anomalies based solely on encoded constellation diagrams.

We also assess the benefit of the proposed approach by measuring the runtime of the anomaly detection approaches. Using the existing dataset, we extracted the average and standard deviations for the normal and anomalous scenarios. These values are applied to parameterize a normal distribution used to synthetically generate another dataset for the runtime assessment. Fig. 4 shows the runtime of the algorithms when the dataset is scaled from 100 to 10,000 samples. For the un-encoded datasets, we were able to process up to 5,000 samples (due to out-of-memory problems). The runtime results show that DBSCAN-AE runs up to 200 times faster than DBSCAN-U if a dataset contains a small number of samples (e.g., 100 samples). For a dataset with many samples (e.g.,

5000), the corresponding speedup is at least 10-fold. Moreover, DBSCAN-AE has a similar runtime to OCSVM. The proposed approach allows for processing up to 2,500 samples within less than one minute, which is a desired monitoring window for optical networks [11]. Assuming a shorter runtime budget for the anomaly detection algorithm of 10 s (e.g., due to other processing required within the monitoring window), it is possible to process up to 1,000 samples.

V. CONCLUSION

The proposed approach sieves relevant features from images of constellation diagrams using a CNNs-based encoder. The extracted features are exploited by an UL clustering algorithm – we use the DBSCAN – to detect anomalies. Performance evaluation under the four considered anomalous conditions indicates that, with a correct configuration, the proposed approach can detect previously unseen optical spectrum anomalies with 100% accuracy, outperforming the OCSVM approach. Moreover, the proposed autoencoder brings up to a 200-fold runtime reduction. The proposed approach paves the way for more interoperable ML methods in optical networks suitable to operate in parallel with other models that exploit constellation diagrams as unified input to assess a broad range of performance metrics. Furthermore, the same CNN can be extended with other outputs to provide several OPM estimations obtained by a single execution/inference.

REFERENCES

[1] C. Natalino *et al.*, "Constellation diagrams for spectrum anomaly detection in optical networks," 2021, DOI: 10.21227/g9s9-ba02.
[2] F. Paolucci *et al.*, "Network telemetry streaming services in sdn-based disaggregated optical networks," *IEEE/OSA J. Lightwave Technol.*, vol. 36, no. 15, pp. 3142–3149, Aug 2018.
[3] D. Rafique and L. Velasco, "Machine learning for network automation: overview, architecture, and applications [invited tutorial]," *IEEE/OSA J. Opt. Commun. Netw.*, vol. 10, no. 10, Oct 2018.
[4] X. Chen *et al.*, "Self-taught anomaly detection with hybrid unsupervised/supervised machine learning in optical networks," *IEEE/OSA J. Lightwave Technol.*, vol. 37, no. 7, pp. 1742–1749, April 2019.
[5] M. Furdek *et al.*, "Experiment-based detection of service disruption attacks in optical networks using data analytics and unsupervised learning," in *Metro and Data Center Optical Networks and Short-Reach Links II*, vol. 10946, San Francisco, CA, USA, 2019.
[6] F. N. Khan *et al.*, "Modulation format identification in coherent receivers using deep machine learning," *IEEE Photonics Technology Letters*, vol. 28, no. 17, pp. 1886–1889, Sep. 2016.
[7] D. Wang *et al.*, "Intelligent constellation diagram analyzer using convolutional neural network-based deep learning," *Opt. Express*, vol. 25, no. 15, pp. 17 150–17 166, Jul 2017.
[8] S. Savian *et al.*, "Joint estimation of IQ phase and gain imbalances using convolutional neural networks on eye diagrams," in *Conference on Lasers and Electro-Optics (CLEO)*, San Jose, CA, USA, May 2018.
[9] S. Peng *et al.*, "Modulation classification based on signal constellation diagrams and deep learning," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 3, pp. 718–727, March 2019.
[10] E. Schubert *et al.*, "DBSCAN revisited, revisited: Why and how you should (still) use DBSCAN," *ACM Trans. Database Syst.*, vol. 42, no. 3, pp. 19:1–19:21, Jul. 2017.
[11] A. P. Vela *et al.*, "Distributing data analytics for efficient multiple traffic anomalies detection," *Computer Communications*, vol. 107, 2017.
[12] S. Varughese *et al.*, "Identification of soft failures in optical links using low complexity anomaly detection," in *OFC*, USA, 2019, p. W2A.46.
[13] B. Schölkopf *et al.*, "Estimating the support of a high-dimensional distribution," *Neural Computation*, vol. 13, no. 7, pp. 1443–1471, 2001.
[14] N. Skorin-Kapov *et al.*, "Physical-layer security in evolving optical networks," *IEEE Commun. Mag.*, vol. 54, no. 8, pp. 110–117, 2016.
[15] VPIphotonics GmbH, "VPIphotonics Design Suite[TM]," www.vpiphotonics.com., accessed 16 September, 2020.