Root Cause Analysis for Autonomous Optical Network Security Management

Carlos Natalino, *Member, IEEE*, Marco Schiano, *Senior Member, IEEE*, Andrea Di Giglio, *Member, IEEE*, Marija Furdek, *Senior Member, IEEE, Optica*

Abstract—The ongoing evolution of optical networks towards autonomous systems supporting high-performance services beyond 5G requires advanced functionalities for automated security management. These functionalities need to support risk reduction, security diagnostics and incident remediation strategies. To cope with evolving security threat scenarios, security diagnostic approaches should be able to detect and identify the nature not only of existing attack techniques, but also those hitherto unknown or insufficiently represented. Machine Learning (ML)-based algorithms have been shown to perform well when identifying known attack types, but cannot guarantee precise identification of unknown attacks. This makes Root Cause Analysis (RCA) a crucial tool to enable timely attack response when human intervention is unavoidable.

We address these challenges by establishing an ML-based framework for security assessment and analyzing RCA alternatives for physical-layer attacks. We first scrutinize different Network Management System (NMS) architectures and the corresponding ML-based security assessment functionalities. We then investigate the applicability of supervised and unsupervised learning (SL and UL) approaches for RCA and propose a novel UL-based RCA algorithm called Distance-Based Root Cause Analysis (DB-RCA). Extensive validation of the framework's applicability and performance in the context of autonomous optical network security management is carried out using an experimental physical-layer security dataset, evaluating the benefits and drawbacks of the SL- and UL-based RCA techniques. Besides confirming that SL-based approaches can be trained to provide precise RCA output for known attack types, the study shows that the proposed UL-based RCA approach offers meaningful insights into the properties of anomalies caused by novel attack types, thus supporting the human security officers in advancing the physical-layer security diagnostics.

Index Terms—Optical networks, physical layer security, supervised learning, unsupervised learning.

I. INTRODUCTION

As the key technology underpinning the global communication network infrastructure, secure and reliable operation of optical networks is a necessary enabler for the network evolution beyond 5G. Due to the ultra-high data rates carried over the optical fiber, even short disruptions at the physical layer can affect a multitude of aggregated upper layer services,

which makes the optical network an enticing target of man-made attacks. The current trends of optical/wireless integration and extending the optical domain towards the network edge in order to reduce latency further aggravate the damaging potential of physical-layer attacks compared to current architectures with optical-electrical-optical conversions at segment exchange points.

Attacks aimed at physical layer disruption may vary in their levels of disturbance, sophistication, or traceability, among other traits. One of the well-known service degradation techniques reported in the literature is jamming [1], where a harmful signal is inserted in the fiber to either add un-filterable noise to a signal at the same wavelength (in-band jamming), or to deprive the co-propagating useful signals of optical amplifier gain (out-of-band jamming). Another technique is a polarization modulation attack, where the fiber is squeezed at a high frequency, resulting in very fast changes of the polarization of light that cause errors when the coherent receiver's polarization recovery algorithm cannot compensate for the fast fluctuations [2].

Coping with evolving security threats in the frame of network transformation into automated high-performance systems requires intelligent approaches for autonomous security management embedded into services' life cycles. Such frameworks should encompass functions for risk mitigation, attack cognition and incident remediation, incorporated into a cognitive NMS (C-NMS) paradigm [3]. *Risk mitigation* can be achieved through attack-aware network design and service provisioning, where spectral and spatial resources are intelligently assigned to limit the potential damage from jamming [4]–[6], or through deployment of devices that can filter out harmful signals [7]. Due to the absence of exact analytical models of physical-layer attack effects, *attack detection* has benefited greatly from ML techniques capable of perceiving the resulting subtle changes in optical signal parameters [2], [8]. An analysis of the role of ML in security management automation can be found in [9]. Fast and accurate detection and identification of attacks is crucial for appropriate and effective *remediation* of security breaches, which should encompass service recovery, threat neutralization and network adaptation. However, due to the constantly evolving threats, ensuring the availability of attack identification for every kind of attack is impractical, and human intervention is unavoidable when the identification of the attack is not possible. To this end, the root cause of a detected anomaly should be investigated before triggering the correct countermeasures.

The first stage of incorporating ML-supported RCA in

Carlos Natalino and Marija Furdek are with the Electrical Engineering Department, Chalmers University of Technology, Gothenburg, Sweden. Marco Schiano and Andrea Di Giglio are with Telecom Italia, Turin, Italy.

The results reported in this paper are partially supported by VR (2019-05008), VINNOVA (AI-NET-PROTECT – 2020-03506), and the European Commission 5GPPP TeraFlow (101015857).

We gratefully acknowledge Infinera for providing the Groove G30 transponders used in the experiments.

Manuscript received XXX XX, 2021; revised XXX XX, 2022.

Security Operations Center (SOC) processes is expected to aid the decision process of a human security operator rather than eliminating the human from the loop [9]. Therefore, it is important to provide insightful information about the detected anomalies. RCA in networks is typically addressed using Supervised Learning (SL) techniques that rely on prior knowledge of anomalies to be detected [10], [11]. However, these approaches may not be applicable to physical-layer security scenarios where a representative labeled dataset covering all attack scenarios may not be available, or new threats may emerge at any time, or models trained for one optical channel may not be applicable to other channels. Therefore, in this work, we establish an ML-based framework for security assessment. Then, we propose an Unsupervised Learning (UL)-based RCA algorithm named DB-RCA that does not require prior knowledge on the anomalies caused by physical-layer attacks. The algorithm, which extends our preliminary study in [12], is validated on an experimental physical-layer security dataset. Our performance assessment includes the RCA outputs of eXtreme Gradient Boosting (XGBoost), a known SL model that enables RCA, in addition to our proposed DB-RCA. By analyzing the output of the SL- and UL-based algorithms we can assert that, although significantly different, the outputs of both approaches can provide meaningful insight when investigating the root cause of an attack.

The contributions of this paper, and extensions with respect to [12] and [9], can be summarized as follows:

- We analyze the different NMS architectures and the corresponding ML-based security assessment functionalities.
- We investigate RCA approaches and establish an RCA framework applicable to the security assessment task.
- We propose a UL-based DB-RCA algorithm, extending the initial study in [12].
- We validate its applicability to autonomous optical network security management using an experimental physical-layer security dataset.
- We scrutinize the performance, benefits and drawbacks of the SL- and UL-based RCA approaches.

The remainder of the paper is organized as follows. Sec. II discusses the NMS progress towards automated security diagnostics enabled by ML for attack detection and RCA. Sec. III introduces the concept of RCA, establishes an RCA framework, examines its importance in the context of autonomous optical network operation, and outlines nuances between SL and UL when performing RCA. Our DB-RCA algorithm is introduced in Sec. IV. The use case, performance assessment of the different RCA approaches and a discussion of the results is presented in Sec. V. Finally, Sec. VI concludes the paper and discusses remaining challenges.

II. PROGRESS TOWARDS AUTOMATED SECURITY DIAGNOSTIC PROCEDURES

The introduction of telemetry systems and advanced ML techniques enables NMSs to evolve towards supporting automated security diagnostics. The evolution of the NMS architecture is illustrated in Fig. 1, starting from a legacy

Fig. 1. Network Management System (NMS) architecture evolution with the introduction of telemetry systems and advanced ML techniques for security diagnostics.

scenario characterized by traditional NMS and describing a possible evolution path towards optical networks with embedded telemetry systems and optical security ML tools.

A. Network Assurance approach based on traditional NMSs

Today, most optical networks are still operated by traditional NMSs, as illustrated by Fig. 1a. They are characterized by collecting Optical Performance Monitoring (OPM) data every 15 minutes with very limited OPM data storage capabilities.

The limitations of those technologies force operators to implement elementary reactive strategies based on alarm monitoring and manual intervention as countermeasures to both faults and attacks. In case of an attack, an alarm is typically raised in the NMS and the operator is alerted in case an OPM value exceeds a predetermined threshold. The NMS itself provides to the operator remedy tables that propose plausible causes for each alarm, suggesting also further checks and appropriate remedy actions.

Due to the limited historical records and the requirement of manual checks, security assessment is very difficult and requires highly-trained and experienced operators that understand the specifics of the system. Moreover, the use of predefined thresholds is unscalable and unreliable, as some attack techniques cause only minor variations in OPM values and each optical channel might need a specific fine tuning of its thresholds [2].

B. NMSs with telemetry functions

A first evolutionary step in optical network management comprises the development of telemetry systems embedded in the NMS, that collect large OPM data records every second (or few seconds) and store those records in a database, as illustrated by Fig. 1b. This technology is already proposed by many network system manufacturers for their next generation

Fig. 2. Workflow of an NMS with ML-based attack detection (and identification) and Root Cause Analysis. The anomaly detection (and identification) (AD/ADI) module outputs the attack class/cluster and known/unknown signature based on the appropriate use of SL and SSL/UL models.

products. Nevertheless, even with this advanced telemetry system, attack diagnostics is similar to the traditional one, but the checks suggested by remedy tables can be done on OPM historical data series rather than on log files with few tens of entries.

In any case, the historical data series analysis, which is performed manually by the operator, is very complex, time consuming, and requires special data processing skills seldom included in the training of Network Operation Center (NOC) operators. Based on their standard fault-oriented mindset, the network assurance operator has no effective tools to analyze what is happening in the network in case of an attack and no criteria for applying specific data analysis tools on the OPM dataset.

C. NMSs with telemetry and attack detection ML functions

A network malfunction is caused by either a failure or a malicious attack. Often, an attack and a failure can have similar *symptoms* but need different *treatments*. For this reason it is very important to have, alongside telemetry systems which continuously monitor the status of the connections and the quality of the optical signals, a system that, based on these data, is able to recognize that the malfunction is caused by a malicious attack and classify such attack based on a previously identified set. Such an NMS architecture, capable of detecting, identifying and locating the source of an attack, is illustrated by Fig. 1c. We have already demonstrated the effectiveness of ML techniques to detect and classify network attacks [2]. When those ML tools are embedded in future generation NMS, they can provide an effective support for the NOC operators: they will gain information on the nature of the network attack directly from the ML system in quasi real time, and they will use their knowledge and experience to apply appropriate countermeasures.

D. NMSs with telemetry, attack detection and Root Cause Analysis ML functions

To support automated security diagnostics, NMS should also have embedded functionalities for tracking newly emerging attacks. As an evolution of the previously described scenario, where the NMS has an embedded algorithm capable of discerning between a failure and an *a priori* known attack and, correspondingly, classifying the attack. Leveraging currently available technologies, properly addressing newly emerging attacks unavoidably requires human intervention. However, a step forward is given by an ML-driven RCA functionality due to its ability to provide an initial insight into the effects of a novel threat due to its capability to provide an initial insight into the effects of a novel threat.

NMS architecture with RCA is illustrated in Fig. 1d. This approach does not rely on prior knowledge of attack consequences, but can, nonetheless, provide meaningful insight into their effects on network performance and aid the operators in determining the most effective security countermeasures.

III. ROOT CAUSE ANALYSIS IN OPTICAL NETWORKS

Root Cause Analysis (RCA) is a term used to define methods that aim at identifying the root causes of faults, problems, or anomalies. Application areas include production engineering and management, IT operations, and telecommunications. Depending on the area where RCA is applied, it might include different steps and techniques. With the advent of the use of ML techniques across a wide range of industries, the black-box nature of many ML models increases the need for RCA specifically tailored for analyzing the output of the models. These ML-based RCA methods analyze the same data that can be input to an ML model, and produce metrics that indicate which aspects or features of the data are more likely to be decisive to the ML output.

In the context of NMSs, RCA can be used to assist the network management staff with the difficult task of analyzing the OPM parameters and identifying the cause of an anomaly. The difficulty of this task is exacerbated when the ML algorithms in place raise an alarm without giving further insight such as identifying the problem.

Fig. 2 illustrates the workflow of an NMS and the role of the network management staff in addressing detected anomalies, which can be further enhanced by anomaly identification. OPM samples are received periodically, which triggers the entire process. An ML-based model performs anomaly detection (and possibly identification). The right-hand side of the figure expands the related procedures. The depicted workflow combines SL and Semi-Supervised Learning (SSL)/UL models to obtain a reliable anomaly assessment. First, an SL model performs multi-class classification, which results in anomaly detection and identification in addition to a confidence score associated with the output of the model. If the confidence level is above a desired level, the known anomaly signature is reported. On the contrary, if the confidence level is below the desired level, the output of the SL model is not considered reliable enough and, instead, anomaly detection is performed by either SSL or UL. In this case, only binary classification or clustering are possible, and the output specifies whether or not an anomaly is detected. If the SSL/UL model detects an anomaly, it reports the detection flagged as an unknown anomaly signature. Note that the described workflow also functions when no SL model is available to detect and identify an anomaly. In this case, the received OPM sample is used by the SSL/UL model to perform binary classification, always flagging the detected anomaly as unknown. Long-term storage is usually needed in this case to store either a set of the most recent OPM samples, or the ML model itself, or both. The output of this module is evaluated, and if no anomaly is detected, the process idles and waits for the next monitoring cycle.

On the other hand, if an anomaly is detected, the process for anomaly identification and mitigation starts. First, if the ML model is capable of anomaly identification, or if a standalone anomaly identification is in place, the NMS checks if the anomaly signature is known. If so, automated anomaly localization and mitigation developed *a priori* are executed.

RCA becomes vital in the case where the anomaly signature is not known. In this case, manual mitigation of the anomaly needs to be tailored by the network management staff. The RCA is helpful at this stage when the network management staff needs to understand what caused the ML model to identify the current network state as anomalous. Without RCA, the network management staff is left with the manual investigation procedures described in Sec. II. RCA provides useful insights, often enriched with graphical representation that can enhance and speed up the analysis phase.

It is important, however, to mention that RCA applied to ML models has some limitations. First and foremost, RCA is based on data, and therefore, can only highlight the possible causes that are represented in the data. This is especially important in optical networks, where minor electrical and temperature variations can affect the correct operation of optical devices

Fig. 3. The operating principle of SL and UL models for training (only in the SL case) and inference. Note that usually $|S| \ll |T|$.

[11], [13]–[16]. These variations are usually not represented in the OPM data collected from the coherent transceivers, which usually only contain optical-related parameters (i.e., do not include electrical current, power, or temperature). Therefore, these electrical and/or temperature variations will manifest in the form of OPM parameter changes. These limitations can be addressed by including, integrating and consolidating further monitoring data (e.g., from the electrical power grid and/or cooling system) into the analysis, which is out of the scope of this work. In the next sections, we introduce RCA techniques for SL and UL.

A. Supervised Learning for Root Cause Analysis

Supervised Learning (SL) models comprise two different stages represented in Fig. 3a. First, the model is trained over a training dataset which needs to be labeled, i.e., for each sample, a label specifies what is expected as an output from SL model. During training, the model abstracts the important information present in the dataset that allows it to adjust and store its internal representation to facilitate the association between each input and the label. For instance, traditional Feed-Forward Neural Networks (FFNNs) represent the information in the form of weights and biases. Decision Trees (DTs), on the other hand, store the information in the form of thresholds and conditional statements. Once trained, SL models can perform inference by analyzing only a single sample (i.e., there is no need to analyze previous samples). This means that SL models are usually resource-intensive during training, but lightweight during inference.

One drawback of SL models, in particular in the context of anomaly detection and security assessment, is that their performance expectations can only be drawn over data that is used for training or test purposes. This means that new anomalies or security threats (previously unseen by a model) can potentially remain undetected for long periods of time unless they are mistakenly classified. In the context of optical networks, this issue becomes more critical as each optical channel may need a specific SL model trained just for that channel, given that, to the best of our knowledge,

the applicability of a single model addressing multiple optical channels established over a different set of optical devices has not been reported in the literature so far. This is due to the fact that different optical channels use different spectrum and traverse different equipment, which differentiates their behavior enough to render an SL model trained for another optical channel unsuitable. Another potential reason is the fact that different optical channels may be designed for different OPM levels such as Optical Signal-to-Noise Ratio (OSNR), modulation format, etc. During operation, this translates either to a longer deployment time for optical channels (due to the need to collect enough samples for every anomaly/attack case), or running the optical channel with user traffic without proper anomaly/attack detection/classification assessment in place [2].

Another drawback of SL models is their trade-off between accuracy and interpretability. Artificial Neural Networks (ANNs) usually achieve the highest accuracy among the SL models when their hyperparameters are appropriately tuned, but are considered a black-box model. This is due to the fact that their internal representation (in the form of weights and biases) is hard to be interpreted. On the other hand, traditional DTs do not usually reach the same level of accuracy as ANNs, but are easily interpreted and can be implemented with simple conditional statements. This makes DTs ideal for use in the RCA for SL models.

The eXtreme Gradient Boosting (XGBoost) algorithm [17] was proposed to mitigate the usual accuracy limitations of DTs, and reduce the trade-off between accuracy and interpretability. XGBoost implements gradient boosting for DTs, and uses an ensemble of DTs to improve its accuracy. The training procedure uses the gradient descent algorithm, which enables the use of different loss functions depending on the task at hand (e.g., cross-entropy for classification tasks and mean squared error for regression). The gradient descent decides whether to add new nodes/leaves to the current DT structure such that the loss is minimized. XGBoost is in the same category as Random Forests (RFs), i.e., ensemble learning algorithms based on decision trees. The main difference between the two is that in RFs each decision tree is independent from another, while in XGBoost the decision tree $n+1$ focuses on improving the accuracy obtained by decision tree n .

Fig. 4 illustrates an example decision tree for a binary classification task in an optical network. Each node in the tree represents a conditional statement, similar to traditional DTs. However, leaves represent the probabilities of a sample belonging to a given class, as opposed to traditional DTs where leaves represent decisions. Note that the example shows a tree for binary classification, so a probability closer to zero means that the sample is likely of the first class, while a probability closer to one means the sample is likely of the second class. By using an ensemble of DTs, XGBoost collects these probabilities across the different DTs to define the output.

XGBoost is particularly useful for RCA due to its importance score that results from the training process. The importance score is attributed to each feature (or OPM parameter in our context) and represents the number of times that the

Fig. 4. An illustrative example of a decision tree for binary classification using OPM parameters. Note that we consider normalized OPM parameter values.

Fig. 5. An illustrative example of feature importance obtained from the XGBoost algorithm.

feature is used to add a new split/node. The intuition is that a tree with an additional split/node is more accurate than one without the split. Therefore, the more times a feature is used in splits, the more important it is for the particular effect under analysis. Fig. 5 illustrates the normalized feature importance resulting from a XGBoost training over a hypothetical dataset collected from optical transceivers for classification of anomalies. We can see that pre-FEC Bit Error Rate (BER-FEC) is the most important feature, with Optical Power Received (OPR) accounting for approximately 30% of the importance of BER-FEC. OSNR, Q-factor and Chromatic Dispersion (CD) have a minor importance. This importance score can be directly linked to the particular effect under analysis and used to tailor mitigation strategies.

B. Unsupervised Learning and Root Cause Analysis

Unsupervised Learning (UL) models are tailored to identify anomalous samples by considering an observation window containing a number of samples. The observation window usually has a few hundreds of samples. The UL models assume that at the beginning of observation, only a few samples in the observation window will contain anomalous behavior. By considering the distance between samples to group them, it is expected that anomalous samples will be far from the samples that characterize normal behavior, and will, therefore, be flagged as anomalies.

Fig. 3b illustrates the operating principle of UL models. As opposed to SL, UL does not require training, which translates into neither using a labeled (training) dataset nor storing of any internal states. This means that it needs to receive an entire observation window every time an inference is needed. Therefore, inference is costlier in UL than in SL models in terms of complexity and runtime [18].

Compared to SL (illustrated in Fig. 3a), UL models have several benefits important for their use in the autonomous operation of optical networks. The absence of training brings two important advantages: (i) these models do not need to be tailored for each optical channel, and a single implementation can be used for all optical channels in a network; and (ii) the models can start their operation shortly after the optical channel is set up, i.e., there is no need to collect a training dataset. Finally, given the general ability to detect anomalies, these models can be potentially applied to any sort of anomaly (e.g., attack, malfunction), unlike SL models that need labeled samples and training for each and every anomalous condition.

On the other hand, there are a few drawbacks of UL models when compared to SL. First, UL models can only detect the anomaly, without any further insight. On the contrary, SL models are able to identify the anomaly as long as they have been trained for it. Moreover, SL models are typically exposed to a larger number of samples during training, which allows them to identify subtle but consistent properties of the system behavior. By analyzing a relatively small number of samples compared to a training dataset for SL models, UL models rely on observing the unstable behavior of the system upon an anomaly occurrence to detect significant deviations from normal operation. As previously mentioned, UL models usually have more complex inference than SL models, although some efforts have been made to combine SL and UL to provide lightweight inference [19]. Finally, UL models usually have a few hyperparameters that define how strictly the relationships among samples will be evaluated, and defining these parameters without knowledge of the anomalies is challenging. One possible way to tune hyperparameters using only normal samples is to start from a more relaxed set of parameters and progressively make them more strict until the algorithm starts presenting false positives. At this point, one can decide on an initially acceptable level of false positives. During network operation, as anomalies are detected and validated through, e.g., an RCA framework, the hyperparameters can be further optimized.

These differences between SL and UL models significantly

Algorithm 1 Distance-based RCA

Data: Set of features F , (pre-processed) dataset X , anomaly detection output Y

Result: Anomaly vector AV

```

1  $N \leftarrow \bigcup_{i=1}^{|X|} \{X_i\} : Y_i \geq 0$ 
2  $A \leftarrow \bigcup_{i=1}^{|X|} \{X_i\} : Y_i = -1$ 
3  $P \leftarrow \text{cluster in } N \text{ closest to } A$ 
4  $AV_i \leftarrow \frac{\sum_{j=0}^{|A|} A_{i,j}}{|A|} - \frac{\sum_{k=0}^{|P|} P_{i,k}}{|P|}, i=0..F$ 
5 return  $AV$ 

```

Fig. 6. An illustrative example of the Anomaly Vector (AV) obtained from the proposed RCA framework.

change the way that RCA is performed in each scenario. With SL models, RCA aims at finding the ground truth importance of features based on extended (and usually far) past behavior (represented by the training dataset). On the other hand, RCA based on UL models aims at finding the importance of features based on a limited number of (close to) real-time observations. In the following section, we describe the algorithm proposed in this paper that aims at providing a viable RCA approach for UL models.

IV. DISTANCE-BASED ROOT CAUSE ANALYSIS (DB-RCA)

This section describes the main contribution of this paper, i.e., the algorithm for UL-based RCA called DB-RCA. The main output of the algorithm is the Anomaly Vector (AV). The AV enriches the anomaly detection algorithm by representing the average distance between anomalous samples and their closest cluster of normal samples. The results of the AV can be graphically shown to the network management staff together with other data relevant for anomaly mitigation. Furthermore, the AV can be integrated into automated anomaly localization and mitigation procedures.

Alg. 1 presents the pseudocode of the proposed algorithm. The algorithm takes as input the set of features F that describes which features are present in the dataset, and the (pre-processed) dataset X containing the features for every sample. Additionally, the algorithm also receives the output from the anomaly detection algorithm Y , where the number of elements in Y is equal to the number of elements in X . Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [20] is one of the UL models that can be used

for anomaly detection, i.e., to produce Y , which is also the approach taken in this work. In any case, the UL algorithm used to detect an anomaly will provide at least two clusters for RCA. For a sample with index i , we assume that $Y_i = -1$ when sample X_i is flagged as anomaly, and $Y_i \geq 0$ when the sample is considered normal, in which case the value of Y_i acts also as the cluster index for the sample indexed by i .

The algorithm begins by obtaining a cluster that contains normal samples (line 1) and anomalous samples (line 2) from X , denoted by N and A , respectively. Then, the algorithm selects the cluster closest to the anomalous samples (line 3) and computes the anomaly vector as the difference between the average values of the features in the two clusters (line 4). By considering only the closest cluster, the algorithm is more likely to observe smaller variations in the features that led the samples to be positioned outside of the cluster area. Moreover, this limits the number of samples whose distance needs to be computed for each feature. The AV is returned (line 5). Note that our proposed algorithm has a small impact on the overall system runtime, since it is only executed when an anomaly is detected. Note also that the algorithm is extensible enough to be included in pipelines using other models. Finally, the algorithm uses the same data used for DBSCAN (or another UL algorithm), in addition to its outputs, without keeping any state between inferences. This means that the algorithm does not change the fundamental properties of DBSCAN and therefore can also be considered a UL approach.

Fig. 6 shows an illustrative example of how the AV would be presented to the network management staff for a hypothetical detected attack. Note that the AV enables a network security specialist to identify the level of deviation from the normal for the considered OPM parameters (e.g., DGD and BER-FEC deviate the most, while OSNR deviates the least in this hypothetical case). It also indicates the increasing or decreasing trends in the parameter variations, represented by positive or negative AV values, respectively.

V. PERFORMANCE ASSESSMENT

This section presents a detailed analysis of the performance of the supervised and unsupervised learning models presented in this work, focusing on their accuracy and RCA properties. The accuracy performance is assessed in terms of the *f1 score*, which balances the importance of false positives and false negatives into a single metric equal to 1 when the model is perfectly accurate in detecting attacks, and smaller otherwise. Then, we analyze the importance metrics output by XGBoost and the proposed UL-based RCA algorithm. Finally, we discuss the characteristics and properties of the evaluated approaches.

A. Use Case

We validate the proposed algorithm on a physical layer security use case where anomalies are characterized by physical layer attacks launched over optical channels. OPM samples are collected from an experimental optical network testbed with coherent transceivers, Reconfigurable Optical Add-Drop

Multiplexers (ROADMs) and Erbium-Doped Fiber Amplifiers (EDFAs). A detailed description of the experiments and testbed can be found in our previous work [2]. Two optical 200 Gbit/s polarization multiplexed 16 quadrature amplitude modulation (16QAM) signals are monitored, characterizing a baseline (no attack) scenario. Then, three attack strategies are launched in the network, namely In-Band (IB) and Out-of-Band (OOB) jamming, and Polarization Modulation (PM). For each attack strategy, a light and a strong attack condition is used, resulting in 7 different attack scenarios: Light In-Band jamming (INBLGT), Strong In-Band jamming (INBSTR), Light Out-of-Band jamming (OBLGT), Strong Out-of-Band jamming (OOSTR), Light Polarization Modulation (POLLGT), and Strong Polarization Modulation (POLSTR). In each minute, and for each attack condition and optical channel, a script collects the following OPM parameters over the course of 24 hours: BER-FEC, post-FEC Bit Error Rate (BER-PF), Loss of Signal (LOS), OPR, CD, Differential Group Delay (DGD), OSNR, Polarization Dependent Loss (PDL) and Q-factor. For some of the OPM parameters, the minimum and maximum values obtained within the monitoring window are also reported. The resulting dataset is the largest optical security-related dataset reported in the literature, composed of 1440 samples for each scenario, with each sample containing 31 features. The collected dataset is pre-processed by removing samples with missing features and by applying z-score standardization. Finally, for the analysis reported in this section, only the nominal values are considered (we do not consider minimum and maximum values reported by the transceivers).

B. RCA Using XGBoost

For the results presented in this section, we used the Python open-source implementation of the XGBoost model. We configured the model for binary classification (which is appropriate for the attack detection task) and to use a single decision tree. The dataset was split into 50% for training and 50% for testing purposes. The model was trained over 7 different cases, one including all attack scenarios labeled as *attack* (i.e., there is no distinction between attacks) and 6 other cases where only one attack scenario was included. All cases included the baseline samples.

Figure 7 shows the f1 score for all the trained models. It is possible to notice that the XGBoost achieves very good performance in all cases, having its lowest score around 0.985 for the Light In-Band jamming (INBLGT) attack, which still represents a very good performance. For all other cases, the f1 score is at least 0.99.

Figures 8 and 9 show the resulting decision tree for classifying IB jamming and PM strong attacks, respectively. Note that the numerical values for the OPM parameters are standardized, and that each leaf represents the probability of a sample being an attack. We can see that for IB strong attacks (Fig. 8), only BER-FEC and OPR are sufficient to perform the attack detection. On the other hand, for PM strong attacks (Fig. 9), BER-PF and OSNR are needed in addition to OPR. These results make it evident that each attack scenario is

Fig. 7. F1 score for the XGBoost models, trained over all attack scenarios (All) and for each one of the attack scenarios.

Fig. 8. Resulting tree for identifying strong in-band jamming attacks.

Fig. 9. Resulting tree for identifying strong polarization attacks.

best represented by a particular set of OPM parameters. More importantly, it shows that collecting a set of OPM parameters as complete as possible is paramount for enabling a reliable and future-proof security assessment. In other words, certain parameters might not be important for the currently known attacks, but might become important to diagnose new attacks

in the future.

Next, we move our attention to the feature importance scores obtained by XGBoost (in our case, features represent the OPM parameters collected from the transceivers). Fig. 10 shows the feature importance for the models trained over all the attack scenarios (Fig. 10a) and over each individual scenario. When detecting all the attack scenarios (Fig. 10a), OPR and BER-FEC are the two most important features. However, PDL and CD, usually disregarded for most of the analysis in the literature, retain a significant importance for the overall attack detection.

For IB and OOB jamming (Figs. 10b and 10c), the two most important OPM parameters are BER-FEC and OPR. However, starting from the third place, things are different. While Q-factor is the third most important OPM parameter for both intensities of OOB jamming, it is only the fifth for IB. For IB, CD is the third most important OPM parameter. Interestingly, for the IB, the importance for the strong attack is concentrated over BER-FEC and OPR, while for the light attack, the importance is spread over 6 different OPM parameters. Moreover, looking at the physical nature of the attack, CD and DGD features should not be representative of the OOB jamming attack where just OSNR, Q-factor and BER-PF are expected to vary, but still have significant importance for the attack detection.

Fig. 10d shows that for PM, the order of importance changes from the light to the strong intensities. BER-FEC and OPR swap places in terms of importance. BER-FEC is the most important for the light attack while OPR is ranked third. For the strong attack, OPR comes first and BER-FEC in third. Finally, BER-PF, which is not ranked for the IB and OOB jamming attacks, is important for PM, and replaces DGD that was important for the two jamming attacks.

C. DB-RCA using DBSCAN

For the results presented in this section, we assume the use of DBSCAN for the anomaly detection task. DBSCAN [20] is an anomaly detection algorithm that clusters samples based on their pair-wise distances. The algorithm has two parameters: ϵ defines the radius around each sample within which other samples are considered *neighbors* (usually considering Euclidean distance); and M defines the minimum number of neighbors that a sample needs to have in order to be considered a normal sample. Any sample whose number of neighbors is less than M is considered an anomaly. We used the Scikit-Learn Python implementation of the DBSCAN algorithm. We fine tuned the model to obtain the best f1 score according to the dataset and the standardization procedures applied by testing the following sets of parameter values. The radius is varied in the set $\epsilon = \{0.1, 0.5, 1, 1.5, 2, 3, 4, 5, 10\}$, while the minimum number of samples is varied in the set $M = \{3, 5, 8, 10, 12, 15, 20\}$. The best accuracy was obtained with $M=15$ and $\epsilon=1.0$, which results in a 0.034 false positive rate and 0.189 false negative rate. In a real deployment, strategies such as Window-based Attack Detection (WAD) [2] can be used to mitigate the false positive and false negative rates of DBSCAN.

Our algorithm was executed using a 10:1.5 proportion between normal (no attack) and anomalous (attack) samples.

Fig. 10. Feature importance for the decision trees trained to identify all attacks, and each one of them individually.

The experiments reported in Fig. 11 are obtained by running the algorithm for each attack scenario, averaged over 50 runs. Since in this work the focus is on obtaining a representative AV, DBSCAN was executed with a relatively large number of samples, i.e., 200 normal (no attack) samples and 30 anomalous (attack) samples, which represents a reasonable trade-off between accuracy and runtime.

Fig. 11 shows the AVs obtained for the different attack scenarios. Note that light and strong attack conditions are shown in the same plot due to space reasons, but in a real-world scenario, a plot similar to the one seen in Fig. 6 is expected to be shown. Out of the three attack strategies, IB jamming (Fig. 11a) is the one that incurs the highest relative change on the OPM parameters. However, PM incurs the changes over the highest number of OPM parameters.

The AVs for IB and OOB jamming (Figs. 11a and 11b), show interesting differences between the two attacks. While for IB the CD shows no variation, OOB incurs a significant change in its value. OPR has a slight variation, while DGD, OSNR and Q-factor are the most affected OPM parameters. BER-FEC, BER-PF, LOS and PDL do not show significant fluctuations for IB and OOB jamming attacks.

The PM attack shows a very different AV. It affects BER-PF in addition to other parameters, and the OSNR variations are both in the positive and negative directions, indicating that

there might be a fluctuation of this OPM parameter throughout an attack.

D. Discussion

This work analyzes two fundamentally different ways to implement RCA while using ML models. One is the XGBoost, which is a SL algorithm that, based on the training dataset, builds decision trees capable of performing regression or classification. The other is DBSCAN, which is an UL algorithm that has no training and is executed over a set of samples every time an anomaly detection needs to be performed.

XGBoost is designed to capture the behavior of a phenomena (optical physical layer attacks in our case) by looking at a (historical) dataset that has enough samples to properly represent how the different features (or OPM parameters in our case) are affected by the phenomena. In our case, we use 50% of our dataset to train XGBoost, which represents 12 hours of OPM monitoring for each attack scenario. However, using current ML models, this training needs to be performed for each optical channel, and having these long monitoring windows prior to the optical channel use (i.e., prior to start sending user traffic) are impractical in real world deployments.

DBSCAN, on the other hand, is designed to identify changes on the feature (or OPM parameters in our case) trends as they happen. This means that it does not need any prior training,

Fig. 11. Anomaly Vector (AV) for each attack scenario.

and can be used shortly after the optical channel establishment. In our case, we used 200 samples, which represents 3 hours and 20 minutes considering a 1-minute monitoring window. However, this window can be shortened by increasing the monitoring frequency at the beginning of the optical channel operation, or by using a lower number of samples as input to DBSCAN at the expense of potentially lower accuracy. Such lower accuracy can be mitigated by, e.g., leveraging WAD [2].

From their very different approaches to training and inference, it can be expected for the RCA results from XGBoost and DBSCAN to diverge to some extent. Fig. 10 shows that

one of the most important features found by XGBoost is BER-FEC, which presents negligible variations in the AVs for all attack scenarios in Fig. 11. On the contrary, OPM parameters such as DGD show a significant variation in the AVs (Fig. 11), but do not figure among the top five most important features in Fig. 10. These divergences are explained by the different time scales that the algorithms have access to. For instance, the variations caused to BER-FEC by attacks can be very small, but consistent. By analyzing a large number of samples, XGBoost can consistently and confidently identify these small variations that are present in the attack samples and leverage this to build the decision trees. Conversely, OPM parameters such as DGD can have inconsistent patterns during an attack, making it harder to consistently and confidently use it to determine whether there is an attack or not. In this case, using BER-FEC to identify a small but consistent difference is more effective than using DGD. This effect can also be observed with respect to OPR, which shows a small variation in the AVs (Fig. 11a), but figures among the top two most important OPM parameters for the IB jamming (Fig. 10b).

However, a few OPM parameters are identified as valuable by both algorithms. For instance, BER-PF shows variations in the AV only for the PM attack in Fig. 11. Likewise, Fig. 10 shows BER-PF among the important features only for the PM attack.

This demonstrates that the DB-RCA algorithm proposed in this work is an effective way to illustrate the important features that were decisive for the samples to be considered anomalies. Moreover, in some cases, it shows similar results as other algorithms that require much more data to be trained. The AV visualization equips the operators with deeper insight into the anomaly structure and eases the physical interpretation of the anomaly thus complementing the ML-assisted RCA tool. AV visualization goes beyond the typical historical data plot that is provided by NMSs today, where simple time series of OPM parameters are presented to the operators. This is especially significant when a new, previously undetected anomaly is analyzed.

VI. CONCLUSIONS

This work investigated ML-based techniques for RCA as an important enabler of autonomous optical network security management. The considered RCA framework may rely on either existing SL-based approaches, such as XGBoost, or on the newly proposed UL-based DB-RCA algorithm, or both. The DB-RCA algorithm also provides a useful graphical representation of a detected anomaly to the network operator in the form of an anomaly vector. An in-depth analysis reveals the performance of the two fundamentally different approaches and uncovers their advantages and drawbacks when applied to the physical-layer security diagnostics. Although differences are observed, we can see that both methods can show significant insights into the properties of the experienced attack. For future work, we plan to investigate the practical implementation aspects of the frameworks in real-world NMSs and evaluate the applicability of the approaches to other anomaly detection use cases.

REFERENCES

- [1] N. Skorin-Kapov, M. Furdek, S. Zsigmond, and L. Wosinska, "Physical-layer security in evolving optical networks," *IEEE Commun. Mag.*, vol. 54, no. 8, pp. 110–117, Aug 2016, DOI: 10.1109/MCOM.2016.7537185.
- [2] M. Furdek, C. Natalino, F. Lipp, D. Hock, A. D. Giglio, and M. Schiano, "Machine learning for optical network security monitoring: A practical perspective," *J. Lightwave Techn.*, vol. 38, no. 11, pp. 2860–2871, 2020, DOI: 10.1109/JLT.2020.2987032.
- [3] D. Rafique *et al.*, "Cognitive assurance architecture for optical network fault management," *J. Lightwave Techn.*, vol. 36, no. 7, pp. 1443–1450, Apr 2018, DOI: 10.1109/JLT.2017.2781540.
- [4] M. Furdek, N. Skorin-Kapov, and M. Grbac, "Attack-aware wavelength assignment for localization of in-band crosstalk attack propagation," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 2, no. 11, pp. 1000–1009, 2010, DOI: 10.1364/JOCN.2.001000.
- [5] J. Zhu and Z. Zhu, "Physical-layer security in MCF-based SDM-EONs: Would crosstalk-aware service provisioning be good enough?" *J. Lightwave Techn.*, vol. 35, no. 22, pp. 4826–4837, 2017, DOI: 10.1109/JLT.2017.2757956.
- [6] G. Savva, K. Manousakis, J. Rak, I. Tomkos, and G. Ellinas, "High-power jamming attack mitigation techniques in spectrally-spatially flexible optical networks," *IEEE Access*, pp. 1–1, 2021, DOI: 10.1109/ACCESS.2021.3058259.
- [7] K. Manousakis and G. Ellinas, "Crosstalk-aware routing spectrum assignment and wss placement in flexible grid optical networks," *Journal of Lightwave Technology*, vol. 35, no. 9, pp. 1477–1489, 2017.
- [8] M. Furdek, C. Natalino, M. Schiano, and A. D. Giglio, "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*. SPIE, 2019, pp. 73 – 82, DOI: 10.1117/12.2509613.
- [9] C. Natalino, A. Di Giglio, M. Schiano, and M. Furdek, "Autonomous security management in optical networks," in *Optical Fiber Communications Conference and Exhibition (OFC)*, 2021, p. Tu1I.1.
- [10] J. M. N. Gonzalez, J. A. Jimenez, J. C. D. Lopez, and H. A. Parada G, "Root cause analysis of network failures using machine learning and summarization techniques," *IEEE Communications Magazine*, vol. 55, no. 9, pp. 126–131, 2017, DOI: 10.1109/MCOM.2017.1700066.
- [11] C. Zhang, D. Wang, C. Song, L. Wang, J. Song, L. Guan, and M. Zhang, "Interpretable learning algorithm based on xgboost for fault prediction in optical network," in *Optical Fiber Communications Conference and Exhibition (OFC)*, 2020, p. Th1F.3.
- [12] C. Natalino, A. Di Giglio, M. Schiano, and M. Furdek, "Root cause analysis for autonomous optical networks: A physical layer security use case," in *European Conference on Optical Communications (ECOC)*, 2020, pp. We2K–1.
- [13] Z. Wang, M. Zhang, D. Wang, C. Song, M. Liu, J. Li, L. Lou, and Z. Liu, "Failure prediction using machine learning and time series in optical network," *Opt. Express*, vol. 25, no. 16, pp. 18 553–18 565, Aug 2017, DOI: 10.1364/OE.25.018553.
- [14] D. Rafique, T. Szyrkowiec, A. Autenrieth, and J.-P. Elbers, "Analytics-driven fault discovery and diagnosis for cognitive root cause analysis," in *Optical Fiber Communication Conference*, 2018, p. W4F.6, DOI: 10.1364/OFC.2018.W4F.6.
- [15] D. Das, M. F. Imteyaz, J. Bapat, and D. Das, "A non-intrusive failure prediction mechanism for deployed optical networks," in *International Conference on COMMunication Systems NETWORKS (COMSNETS)*, 2021, pp. 24–28, DOI: 10.1109/COMSNETS51098.2021.9352868.
- [16] C. Natalino, A. Udalcovs, L. Wosinska, O. Ozolins, and M. Furdek, "Spectrum anomaly detection for optical network monitoring using deep unsupervised learning," *IEEE Communications Letters*, vol. 25, no. 5, pp. 1583–1586, 2021, DOI: 10.1109/LCOMM.2021.3055064.
- [17] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, NY, USA, 2016, pp. 785–794, DOI: 10.1145/2939672.2939785.
- [18] M. Furdek, C. Natalino, A. Di Giglio, and M. Schiano, "Optical network security management: requirements, architecture, and efficient machine learning models for detection of evolving threats [invited]," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 13, no. 2, pp. A144–A155, 2021, DOI: 10.1364/JOCN.402884.
- [19] X. Chen, B. Li, R. Proietti, Z. Zhu, and S. J. B. Yoo, "Self-taught anomaly detection with hybrid unsupervised/supervised machine learning in optical networks," *J. Lightwave Techn.*, vol. 37, no. 7, pp. 1742–1749, April 2019, DOI: 10.1109/JLT.2019.2902487.
- [20] E. Schubert, J. Sander, M. Ester, H. P. Kriegel, and X. Xu, "DBSCAN revisited, revisited: Why and how you should (still) use DBSCAN," *ACM Trans. Database Syst.*, vol. 42, no. 3, Jul. 2017, DOI: 10.1145/3068335.