Fault Management Based on Machine Learning [Invited]

Luis Velasco and Danish Rafique

Optical Communications Group (GCO), Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
ADVA Optical Networking SE, Munich, Germany
e-mail: lvelasco@ac.upc.edu

Abstract: Machine Learning (ML) brings many benefits for network operation. In this paper, basic ML concepts and its integration into existing network control and management planes are reviewed. Case studies covering fault management are illustrated.

1. Introduction

Transport networks are complex interacting systems, involving cloud operations and core and metro transport. With localized and highly engineered operational tools, it is typical of these networks to take several weeks to months for any changes, upgrades or service deployments to take effect. Machine Learning (ML) is highly suitable for complex system representation as it enables this learning paradigm [1]. ML may be used to achieve network-domain goals including root-cause analysis [2], [3] and failure localization [4], [5], as well as other related capital operational expenditure savings. ML algorithms comprise of a unique ability to learn system behavior from past data, and estimate future response based on the learned system model.

Supported by the recent improvements in computational hardware and parallel computing, the commercialization of big data storage, and processing frameworks, and the introduction of Software-defined networking (SDN) / Network Function Virtualization (NFV) platforms, several networking challenges may be partially or fully addressed using ML paradigms. In this paper, we review several machine learning concepts, tailored to the optical networking industry, and its integration into existing network control and management tools. Finally, we focus on fault management and cover degradation detection and localization, as early detection of equipment failure states and consequent remedial actions can prevent network downtime and enable scheduled preventive maintenance.

2. Overview of Machine Learning

ML is typically thought of as a universal toolbox, ready to be used for *classification*, i.e., identifying to which of a set of categories a new observation belongs to, and *regression*, i.e., estimating the relationships among variables. In fact, it is a diverse field comprising of various constituents and includes data collection and transformation, model selection and optimization, performance evaluation, visualization, integration, etc.

ML approaches may be categorized based on objectives of the learning task, where these objectives may target pattern identification for classification and prediction, learning for action, or inductive learning methods. The algorithms may be further classified into three distinct learning families: *i*) *supervised learning*, *ii*) *unsupervised learning* and *iii*) *reinforcement learning*. *Semi-supervised learning* -or hybrid learning- is sometimes considered as a fourth family, borrowing features from supervised and unsupervised ones.

Let us now introduce the ML families, as depicted in Fig. 1, together with some typical ML algorithms, and their respective applications in optical networking.

2.1 Supervised Learning

Supervised learning makes use of known output feature(s), named *labels*, to derive a computational relationship between input and output data. An algorithm iteratively constructs a ML model by updating its weights, based on the mapping of a set of inputs to their corresponding output features. Examples of algorithms are Artificial Neural Networks (ANN), K-nearest Neighbors, and Support Vector Machine (SVM). Optical networking applications include resource optimization by estimation, and eventual prediction, of network state parameters for a given set of configurations (e.g., symbol rates, optimum launch power, etc.) Another application is ML-driven fault identification, based on historical traffic or network function patterns.

2.2 Unsupervised Learning

While supervised learning provides a clean-slate approach to ML model construction, in practice, labeled data is neither easily accessible nor abundantly available. Unsupervised learning aims to build representation of a given data-set without any label-driven feedback mechanism. Examples of unsupervised learning algorithms are K-mean Clustering, Principal Component Analysis (PCA), and Self Organizing Maps (SOM). Unsupervised learning models may be naturally used for clustering of transport channels, nodes or devices, based on their temporal and spatial similarities. Applications include traffic migration, spectral slot identification, etc.

2.3 Reinforcement Learning

Reinforcement learning refers to ML mechanisms without an explicit training phase. Reinforcement learning

---

Machine Learning

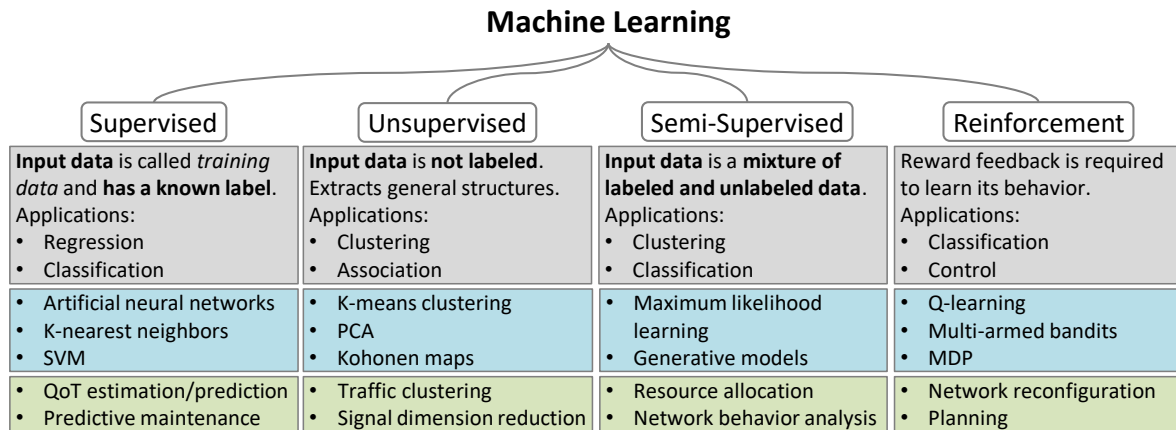| Supervised	Unsupervised	Semi-Supervised	Reinforcement
Input data is called *training data* and **has a known label**. Extracts general structures. Applications: • Regression • Classification	**Input data** is **not labeled**. Extracts general structures. Applications: • Clustering • Association	**Input data** is a **mixture of labeled and unlabeled data**. Applications: • Clustering • Classification	Reward feedback is required to learn its behavior. Applications: • Classification • Control
• Artificial neural networks • K-nearest neighbors • SVM	• K-means clustering • PCA • Kohonen maps	• Maximum likelihood learning • Generative models	• Q-learning • Multi-armed bandits • MDP
• QoT estimation/prediction • Predictive maintenance	• Traffic clustering • Signal dimension reduction	• Resource allocation • Network behavior analysis	• Network reconfiguration • Planning

Fig. 1. ML families.

aims to build and update a ML model based on an agent's interaction with its own environment. The key difference with respect to supervised learning techniques is that labeled input-output features are not provided, but the relationship is rather learned via application of initial model to test data. Examples of algorithms include Q-learning, Random Forest, and Bayesian Network. One of their core applications in optical networks is network self-configuration, including resource allocation and service (re)configurations both for physical and virtual infrastructure.

3. Control, Orchestration and Management

In order to make use of advanced ML models, ML models need to be integrated into existing network software stack. However, multi-layer and multi-vendor network control and management is a complex task in itself, involving core and metro transport, as well as cloud operations. Although SDN has brought agility, flexibility, and scalability, into the network control as compared to traditional control and management platforms, enabling centralized, programmable, and automated services across multiple domains, to attain true network automation, centralized SDN control needs to be augmented with instantaneous data-driven decision-making using advanced monitoring and ML tools, feeding management and control plane alike.

The discussions around SDN have mostly focused on separation of data and control planes, with little attention on operational feedback loop, including monitoring, intelligence and management functionalities. Fig. 2 captures this theme and presents a high-level network architecture, where central offices (CO) consist of intra- and inter- data center infrastructure. The intra-data center resources comprise storage, compute and network, whereas inter-data center connectivity is provided by a transport network. Resources are continuously monitored, exposing real-time network states to the analytics stage, which in turn feeds into the control, orchestration and management (COM) system [6].

This holistic platform not only caters for centralized and programmable control, but also makes ML-driven decisions to trigger actions, essentially connecting data-driven automation with policy-based orchestration and management. To this end, the COM architecture includes the NFV Orchestrator providing network services, the virtual infrastructure manager (VIM) coordinating and automating data center workflows, the network orchestrator adopting hierarchical control architectures with a parent SDN controller abstracting the underlying complexity, and a monitoring and data analytics (MDA) controller that collates monitoring data from network, cloud and applications and contains ML algorithms.
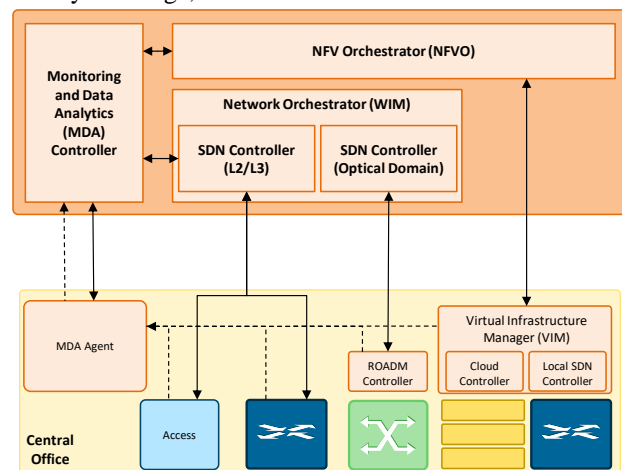

Fig. 2. Autonomic networking architecture.

4. Fault Management

Let us now introduce several applications of ML related to fault management in optical networks. Many commercial equipment tolerates some errors until automatically tearing down the connection when some system thresholds are exceeded. While a restoration procedure could be initiated to recover the affected traffic, it would be desirable to anticipate such degradations and localize the root-cause of the (soft) failure so the lightpath can be re-routed before it is disrupted; note that failure localization is required to exclude the failed resources from path computation, as well as to schedule maintenance tasks. In addition, proactive failure detection would also allow time to plan the re-routing procedure, e.g., during off-peak hours.
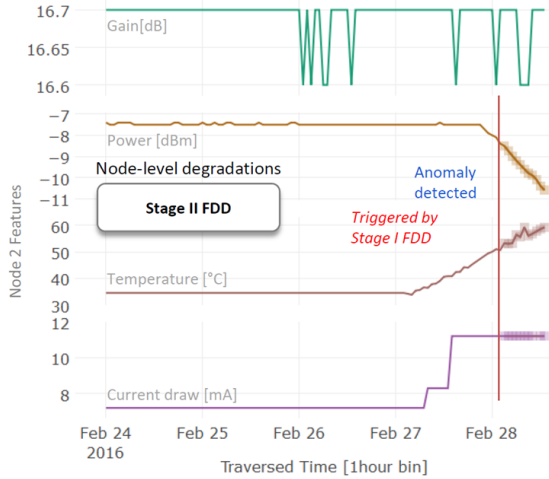
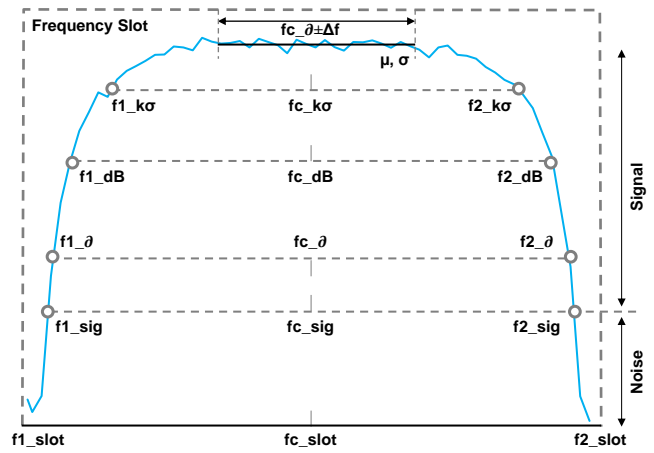Fig. 3. Localized fault discovery and root-cause analysis

Fig. 4. Example of optical spectrum and signal features.

4.1 Predictive Maintenance

In this use case, we address failure detection and localization, and demonstrate an analytics-enabled fault discovery and diagnosis (FDD) architecture, capable of proactively detecting and localizing potential faults and determining the likely root-cause. We carried out an experiment on a network segment consisting of ADVA's FSP 3000 modules, carrying a monitored 100Gb/s transport service. Besides optical power levels, a subscription-based monitoring system measures amplifier gain, shelf temperature, current draw and internal optical power. A fault discovery and diagnosis framework running in the SDN controller is in charge of triggering distributed analysis.

After initial data acquisition, the algorithm is executed in two phases. Firstly, optical power records are partitioned, the *generalized extreme studentized deviate* test is used to detect outliers, and a ANN is used for true fault detection. The second-level analyzes then temperature, amplifier gain, intermediate stage power, and current draw profiles (Fig. 3), where the associated changes in the feature indicate potential root cause. In the case of a high correlation is found between some those features the root-cause of the failure has been localized.

4.2 Optical Spectrum Analysis

Soft-failures can degrade lightpaths' quality of transmission and introduce errors in the optical layer that might impact on the quality of the services deployed on top of such networks. Some soft-failures affect the shape of optical signals and they can be detected by the MDA agents at intermediated nodes analyzing the optical spectrum acquired by local Optical Spectrum Analyzers (OSA). Note that the acquired optical spectra entails large amount of data (e.g., 6,400 frequency-power ($<f, p>$) pairs for the C-band for OSAs with 625 MHz resolution), so local analysis carried out at the MDA agents greatly reduces the amount of data to be conveyed to the MDA controller. Upon detection of a soft-failure, the MDA agent notifies the MDA controller, which is able to correlate notifications received from several MDA agents and for several lightpaths to localize the element that is causing the failure.

Fig. 4 shows an example of the optical spectrum of a 100Gb/s DP-QPSK modulated signal. By inspection, we can observe that a signal is properly configured. However, when a filter failure occurs, the spectrum is distorted, e.g., the optical spectrum can be asymmetrical as a result of one or more filters are misaligned, or the edges of the optical spectrum look excessively rounded as a consequence of the bandwidth of a filter is narrower than the frequency slot width allocated for the signal. Classifiers, based on e.g., SVMs can be used to detect such filter failures in intermediate nodes, so the optical node responsible for the failure can be determined. Once the failure has been localized, the SDN controller can re-route the affected lightpath excluding the failed resource.

5. Conclusions

Traditional networks suffer from largely static operational and optimization practices that limit their scalability and efficiency. ML provides a collection of techniques to fundamentally adapt to the dynamic network behavior. While the application of ML for optical networks is still in its infancy; these learning-based techniques provide a promising platform for end-to-end network automation, including fault management.

References

[1] D. Rafique and L. Velasco, "Machine Learning for Optical Network Automation: Overview, Architecture and Applications," IEEE/OSA JOCN, 2018.
[2] D. Rafique *et al*., "TSDN-enabled network assurance: A cognitive fault detection architecture," in ECOC, 2017.
[3] D. Rafique *et al*., "Analytics driven fault discovery and diagnosis for cognitive root cause analysis," in OFC, 2018.
[4] A. P. Vela *et al*., "Soft failure localization during commissioning testing and lightpath operation," IEEE/OSA JOCN 2018.
[5] B. Shariati *et al*., "Learning from the Optical Spectrum: Failure Detection and Identification [Invited]," IEEE/OSA JLT, 2018.
[6] Ll. Gifre *et al*., "Autonomic Disaggregated Multilayer Networking," IEEE/OSA JOCN, 2018.