Summary

Fiber-optic networks form the backbone of modern telecommunication systems, enabling high-speed data transmission worldwide. Effective monitoring of these networks is essential to identify faults and prevent data loss. However, two major challenges arise in this context: *domain shifts* and *limited supervision.*

Domain shifts occur when the training data differs from those encountered during inference. These discrepancies limit the portability of monitoring solutions, often requiring models to be retrained for each new device. Additionally, limited supervision stems from the high costs and specialized expertise needed to collect annotated datasets, particularly when tailored to specific devices.

Our contribution to improving monitoring in fiber-optic networks is threefold. We first focus on monitoring *Performance Measures* (PMs), which reflect the transmission status at network nodes. Due to the non-linear effects of the optical components, PMs may exhibit significant statistical changes, even when the system is in the *in-control* state, i.e., operate under normal conditions. The *in-control* state is defined solely by a training set, and any deviation not represented in this set is considered an *out-of-control* condition, which must be detected as it indicates a potential fault in the system. Our proposed framework, termed *Hierarchical Change and Anomaly Detection* (HiCAD), effectively distinguishes between these states without relying on device-specific expert rules. It identifies candidate changes, extracts descriptors to capture their underlying patterns, and evaluates whether the detected changes remain within the *in-control* state. If not, an alarm is triggered to indicate a potential issue.

Our second contribution focuses on detecting faults in fiber links, which is critical in large networks and developing regions where breakages occur more frequently. *Optical Time-Domain Reflectometer* (OTDR) devices capture light propagation through the fiber and generate traces that reveal issues, such as faulty connectors or broken fibers, in the form of optical *events*. While manual inspection of these traces is possible, automated methods are preferred for efficiency. Therefore, we developed a 1D *Event-Detection Network* that learns end-to-end to detect events of known classes. Moreover, we exploit an open-set classifier to recognize *unknown events*, i.e., categories of events not included in the training set. Additionally, we address the issue of domain shifts among traces from different OTDR devices. In particular, we designed an *Event-Detection System* that enables models trained on one device to be effectively applied to others by decoupling *localization* from *classification*, using signal processing techniques for event localization and a data-driven model for event classification.

Our third contribution focuses on monitoring optical spectra acquired by *Optical Channel Monitoring* (OCM) modules. Spectra provide a comprehensive view of the signal's power distribution across wavelengths, helping to identify faults that appear in the form of anomalies, i.e., channels that fail to reach their target power levels. During transmission, the non-linear effects of amplifiers and fiber spans introduce distortions that significantly affect the transmitted spectra, making it challenging to detect anomalies. Since distortions occur consistently across the entire spectrum, our method leverages the similarity of the spectrum trends to robustly estimate the expected power of the channels for each frequency location. Specifically, we identify candidate channels using robust fitting techniques and promote trend similarity through our proposed *joint optimization* procedure. Anomalies are detected as deviations from the trend of the channels.

Our advanced monitoring methods effectively learn from limited data and generalize across domains, adapting to various acquisition devices.