

Machine Learning-Enabled Fault Detection in High-Speed Fiber-Optic Networks

Jyrki Aro

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.

jyrkiaro601@unr.edu

Torquato Majewski

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.

tmajewski@ku.edu

Merche Carbajo

School of Computing, Clemson University, Clemson, SC, USA.

merchecarbajo111@clemson.edu

Abstract

High-speed fiber-optic networks constitute the backbone of contemporary digital society, supporting cloud computing, financial settlement, scientific collaboration, mobile backhaul, and increasingly time-critical industrial applications. As transmission rates rise and optical infrastructure becomes more complex, fault detection can no longer rely solely on threshold alarms, periodic testing, or operator experience. Machine learning offers a pathway toward earlier, more granular, and more context-aware identification of degradation, misconfiguration, and physical-layer anomalies. This paper examines machine learning-enabled fault detection in high-speed fiber-optic networks from a system-level perspective. It analyzes the architectural conditions that shape optical faults, the telemetry and sensing infrastructure required for data-driven monitoring, and the trade-offs among detection accuracy, inference latency, interpretability, deployment cost, and operational trust. The discussion emphasizes that fault detection is not merely a model-selection problem but a socio-technical challenge involving governance, workforce capacity, vendor interoperability, sustainability, and equitable access to resilient infrastructure. Through cross-domain comparisons with cloud anomaly detection, power-grid monitoring, and industrial predictive maintenance, the paper argues that successful deployment depends on aligning machine learning pipelines with optical-layer realities, institutional incentives, and public-interest obligations. It concludes with a forward-looking research agenda focused on robust multimodal sensing, uncertainty-aware diagnosis, human-in-the-loop operations, and policy frameworks that treat network resilience as critical infrastructure.

Keywords

machine learning; fault detection; fiber-optic networks; optical telemetry; network resilience; interpretability; infrastructure governance; sustainability.

1. Introduction

High-speed fiber-optic networks have become foundational infrastructures rather than neutral carriers of data. They enable remote work, distributed artificial intelligence, real-time

financial transactions, telemedicine, public safety communications, and scientific instruments that generate enormous data volumes. The continuing expansion of optical capacity has been driven by coherent transmission, dense wavelength division multiplexing, advanced modulation, and increasingly automated network control [1]. At the same time, demand for bandwidth, low latency, and deterministic performance continues to grow across consumer, enterprise, and governmental domains [2]. This growth creates a paradox: optical networks are more capable than ever, yet they are also more difficult to observe, reason about, and repair because their performance depends on tightly coupled photonic, electronic, software, and operational factors.

Fault detection in such environments involves identifying events or conditions that degrade service, threaten availability, or increase the probability of future failure. Traditional approaches often depend on alarm thresholds, optical time-domain reflectometry, performance monitoring counters, and manual correlation across layers. These methods remain essential, but they struggle with gradual degradation, intermittent faults, complex nonlinear interactions, and failures that emerge from configuration drift rather than hardware breakage. Machine learning has therefore attracted attention because it can detect subtle patterns across high-dimensional telemetry, adapt to changing baselines, and support diagnosis before conventional alarms are triggered. However, the transition from promising models to dependable operational systems raises questions about architecture, data quality, governance, and the distribution of responsibility between humans and automated systems.

This paper argues that machine learning-enabled fault detection in high-speed fiber-optic networks should be understood as a system-level design problem. The relevant unit of analysis is not a single classifier but an integrated socio-technical arrangement involving optical hardware, telemetry pipelines, data labeling, model lifecycle management, network operations centers, vendor contracts, regulatory obligations, and sustainability constraints. The sections that follow examine fault dynamics, sensing infrastructure, model families, structural trade-offs, deployment governance, sustainability and policy, case illustrations, and future research directions. The central claim is that technical accuracy alone is insufficient; resilient fault detection requires architectures and institutions that make uncertainty visible, support human judgment, and distribute the benefits of reliability fairly.

2. System Architecture and Fault Dynamics in High-Speed Fiber-Optic Networks

High-speed fiber-optic networks are layered and heterogeneous systems. At the physical layer, optical fibers, amplifiers, multiplexers, transceivers, switches, and reconfigurable optical add-drop multiplexers interact under strict spectral and power constraints. Dense wavelength division multiplexing increases capacity by transmitting many channels across a shared fiber, but it also creates crosstalk, filtering, and power-balancing challenges. International standards define spectral grids and channel spacing, yet real deployments must accommodate legacy equipment, mixed vendor domains, and dynamic reconfiguration [3]. The result is a complex operational environment in which a small physical disturbance can propagate across many logical services.

Fault dynamics in optical networks differ from those in purely electronic systems. Nonlinear propagation effects, amplified spontaneous emission, polarization-dependent loss, chromatic dispersion, and fiber nonlinearity can degrade signals in ways that are difficult to attribute to a single cause [4]. Advanced modulation formats increase spectral efficiency but reduce tolerance to noise and nonlinearity [5]. Analytical models of nonlinear propagation provide important insight, but they remain approximations under real-world conditions involving

mixed fiber types, aging components, temperature variation, and imperfect amplifier control [6]. Consequently, fault detection must often operate in a regime of partial observability, where the observable symptom is separated from the underlying cause by multiple layers of abstraction.

The architecture of wavelength-routed networks further complicates fault localization. A service may traverse multiple fiber spans, amplifiers, ROADMs, and switching nodes owned by different operators. Protection and restoration mechanisms can reroute traffic automatically, masking the original fault while shifting congestion or latency to other paths [8]. In transparent optical networks, the lack of electrical regeneration at every node reduces cost and latency but limits the points at which performance can be measured [9]. Machine learning systems therefore need access to cross-layer telemetry and topology-aware features; otherwise, they may learn correlations that reflect routing decisions rather than physical degradation. The architectural challenge is to create observability without disrupting the very performance characteristics that make high-speed optical transport valuable.

3. Data Acquisition, Telemetry, and Sensing Infrastructure

Machine learning for fault detection depends on data infrastructure that is often less mature than the models themselves. Optical networks generate telemetry from transceivers, amplifiers, optical channel monitors, power meters, temperature sensors, and control-plane events. These sources vary in sampling rate, accuracy, calibration, and temporal alignment. Some metrics are available in real time, while others are collected periodically or only during maintenance. The resulting data landscape is multimodal, sparse in labels, and uneven across vendor domains. Without deliberate instrumentation design, machine learning models may inherit blind spots that mirror existing operational assumptions.

Telemetry also has energy and bandwidth costs. Network power consumption has become a significant concern as traffic grows and as operators seek to reduce operational expenses and environmental impact [7]. Sending all raw optical data to centralized analytics may be impractical, especially when monitoring thousands of channels at high sampling rates. Edge or in-network preprocessing can reduce data volume, but it introduces questions about model placement, privacy, and vendor lock-in. For example, a transceiver may perform local anomaly scoring without exposing raw signal traces, which protects intellectual property but limits root-cause analysis. The design of sensing infrastructure is therefore a governance decision as much as a technical one.

Label scarcity is another structural constraint. Faults are rare by design, and many severe failures are prevented before they recur. Supervised learning requires examples of known faults, yet the most valuable cases may be novel or ambiguous. Semi-supervised, self-supervised, and anomaly-detection approaches can help, but they shift the problem toward defining normal behavior in a constantly changing network. Optical performance monitoring and modulation format identification have advanced through machine learning, yet operational deployment still requires careful validation across equipment generations and environmental conditions [13]. A robust telemetry strategy includes data provenance, versioning, calibration metadata, and retention policies that allow models to be audited and retrained as the infrastructure evolves.

4. Machine Learning Models for Fault Detection and Diagnosis

A wide range of machine learning methods has been applied to optical network monitoring. Classical approaches such as random forests, support vector machines, and gradient boosting

remain attractive because they perform well on tabular telemetry and offer relative interpretability [17]. They can combine optical power readings, error counters, temperature, and configuration data into fault classifiers or remaining-useful-life estimators. Deep learning extends this capacity by learning representations directly from time series, spectrograms, and raw signal samples [18]. Recurrent networks, convolutional networks, and transformer-based architectures have shown promise in time-series classification and anomaly detection [10]. In optical networks, these methods can capture temporal dependencies and nonlinear interactions that simpler models miss [11].

Despite this progress, model choice must reflect operational goals. Fault detection may require low-latency inference at the edge, while diagnosis may tolerate higher latency but demand greater explanatory power. Some tasks, such as quality-of-transmission prediction, benefit from supervised regression using network topology and signal features [14]. Others, such as intrusion or degradation detection, require unsupervised methods because labels are unavailable or unreliable. The literature on machine learning in optical networks emphasizes the importance of feature engineering, data augmentation, and domain adaptation, especially when models trained in one network segment are deployed in another [12]. A model that performs well in simulation may fail in production because the simulation omits amplifier aging, connector contamination, or operational workarounds.

Diagnosis is more difficult than detection. Identifying that a channel is degrading is useful, but identifying whether the cause is a failing amplifier, a misconfigured ROADM, a fiber bend, or a software update requires causal reasoning and contextual knowledge. Machine learning can assist by ranking probable causes, but it rarely eliminates the need for human expertise. Hybrid systems that combine physics-based models with data-driven inference are often more robust because they constrain predictions with known optical relationships. In this sense, the most effective architectures are not pure learning systems but integrated diagnostic platforms that combine signal processing, topology awareness, rule-based logic, and human feedback loops.

5. Structural Trade-Offs: Accuracy, Latency, Interpretability, and Cost

The deployment of machine learning in high-speed optical networks involves unavoidable trade-offs. Higher accuracy often requires more complex models, more data, and more computation. Lower latency requires inference close to the data source, which may limit model size or require hardware acceleration. Interpretability can conflict with accuracy when complex ensembles or deep networks are used, yet operators need to understand why an alarm was raised before taking disruptive action [19]. Cost considerations include not only model training and inference but also sensor deployment, data storage, integration with operations support systems, and ongoing maintenance. A technically superior model may be economically irrational if it requires replacing otherwise functional equipment.

There is also a trade-off between sensitivity and false alarms. In critical infrastructure, false positives can erode trust and cause unnecessary maintenance, while false negatives can lead to outages and safety consequences. The optimal balance depends on the service-level agreement, the cost of downtime, the availability of spare parts, and the regulatory context. For example, a financial trading network may prioritize ultra-low latency and rapid fault detection, whereas a rural broadband network may tolerate longer restoration times if it reduces cost. Machine learning systems should therefore be configurable rather than universal. They must expose uncertainty, support threshold tuning, and allow operators to incorporate local knowledge about risk tolerance.

Another structural tension concerns centralization versus federation. Centralized analytics can aggregate data across domains and improve model performance, but it raises concerns about data ownership, competition, and security. Federated learning offers an alternative by training models across operators without sharing raw data, but it introduces coordination overhead and may struggle with heterogeneous data distributions. In optical networks, where vendors guard telemetry formats and operators compete on service quality, governance arrangements are as important as algorithms. The choice between centralized, federated, and edge-based architectures shapes who can detect faults, who can diagnose them, and who bears responsibility when automated decisions fail.

6. Deployment and Operations: Governance, Robustness, and Human Factors

Operational deployment is where machine learning systems meet organizational reality. Network operations centers already use runbooks, escalation procedures, and change-management processes. A new fault-detection model must integrate with these workflows rather than run parallel to them. If the model produces alerts that are not actionable, operators may ignore them. If it produces recommendations that conflict with vendor procedures, it may be overridden. Successful deployment therefore requires co-design with operators, clear ownership of model updates, and mechanisms for feedback when predictions are wrong. The model is not simply a tool; it becomes part of the sociotechnical system that produces network reliability.

Robustness is a central governance concern. Machine learning models can fail under distribution shift, adversarial manipulation, sensor drift, or software changes. In optical networks, a firmware update may alter telemetry semantics, causing a previously accurate model to degrade silently. Governance must include monitoring for model performance decay, version control for data and models, and rollback procedures when automated actions cause harm. Because optical networks are critical infrastructure, decisions about automation levels should be explicit. Some tasks may be fully automated, such as raising a low-risk anomaly alert, while others require human authorization, such as rerouting protected traffic or dispatching field crews.

Human factors also shape fairness and accountability. If operators do not understand a model, they may either over-trust it or dismiss it. Interpretability is not merely a technical preference; it is a condition for responsible automation. Research on interpretable machine learning argues that high-stakes decisions should favor models whose reasoning can be examined, even at some cost in predictive performance [19]. In optical operations, interpretability can take the form of feature attribution, counterfactual explanations, or topology-aware visualizations. The broader point is that accountability requires traceability. When an automated fault detector contributes to a service outage, regulators, customers, and internal auditors need to reconstruct what happened and why.

7. Sustainability, Equity, and Policy Implications

Fiber-optic networks are often discussed in terms of speed and capacity, but they also have environmental and social dimensions. Manufacturing, deploying, powering, and retiring network equipment consume energy and materials. Machine learning adds its own computational footprint through training and inference, although the incremental cost may be small compared with the energy savings from preventing outages and optimizing routing. Still, sustainability requires life-cycle thinking. A fault-detection system that extends equipment life, reduces truck rolls, and prevents unnecessary replacements can contribute to

environmental goals, while a poorly designed system that generates constant alerts and overprovisioning can increase waste [7].

Equity and policy implications are equally significant. High-speed fiber is unevenly distributed across regions, income levels, and institutional capacities. If advanced fault detection is adopted mainly by large carriers in dense markets, rural and community networks may fall further behind in reliability and cost. Public policy can influence this trajectory through funding for open telemetry standards, shared testbeds, and workforce training. Regulatory frameworks may also need to address transparency, outage reporting, and the use of automated decision systems in critical infrastructure. The question is not only whether machine learning can detect faults, but who benefits from faster detection and who bears the cost of continued fragility.

Policy also intersects with competition and interoperability. If vendors design closed telemetry ecosystems, machine learning models may become locked to specific equipment, reducing bargaining power for operators and limiting innovation. Standards for data models, application programming interfaces, and metadata can enable multi-vendor analytics while preserving security and intellectual property. Governments and industry consortia can support such standards as part of broader digital infrastructure policy. In this sense, fault detection is connected to industrial policy, because the ability to observe and maintain networks is a dimension of national and regional resilience.

8. Case Illustrations and Cross-Domain Comparisons

Useful lessons can be drawn from adjacent domains. Cloud providers use machine learning for anomaly detection across servers, storage, and network flows, often combining unsupervised methods with rule-based guardrails. Power grids use phasor measurement units and predictive analytics to detect instability, but they also maintain strict reliability standards and operator oversight. Industrial predictive maintenance uses vibration, temperature, and acoustic data to anticipate equipment failure, demonstrating the value of multimodal sensing. Optical networks share features with all three: they are critical, geographically distributed, sensor-rich, and dependent on human expertise. The most transferable lesson is that detection systems succeed when they are embedded in maintenance and governance processes, not when they are treated as standalone software.

A specific optical example concerns low-latency, guaranteed-bandwidth services. Recent work on frequency-referenced multiplexing illustrates how physical-layer innovation can provide deterministic performance, which in turn changes fault-detection requirements because latency and bandwidth guarantees become explicit operational commitments [15]. When a network promises strict performance, degradation must be detected before service-level objectives are violated. This places pressure on telemetry granularity, model latency, and control-plane responsiveness. It also raises questions about how guarantees are verified and enforced across multiple operators. The broader implication is that as optical networks become more programmable and service-aware, fault detection must become more tightly coupled with service management and policy enforcement.

Cross-domain comparison also reveals differences. Cloud systems can often migrate workloads away from failing hardware, while optical networks may have fewer redundant paths in rural or undersea segments. Power grids have well-established regulatory regimes for reliability, while optical networks are governed by a mix of telecommunications law, competition policy, and private contracts. Industrial maintenance often has physical access to

failing equipment, whereas optical faults may occur in remote or inaccessible locations. These differences mean that simply importing machine learning methods from other domains is insufficient. Models must be adapted to optical physics, fiber topology, maintenance logistics, and the economics of network operations.

9. Future Directions and Research Agenda

Several research directions can strengthen machine learning-enabled fault detection in high-speed optical networks. First, multimodal sensing should combine optical performance data, environmental measurements, control-plane logs, and maintenance records to improve diagnosis. Second, uncertainty-aware machine learning should become standard, so that operators know when a model is extrapolating beyond its training distribution. Third, physics-informed learning can embed known optical relationships into data-driven models, improving robustness and reducing data requirements. Fourth, federated and privacy-preserving analytics can allow multiple operators to benefit from shared learning without exposing sensitive telemetry. Fifth, human-in-the-loop systems should be designed to capture operator feedback and convert it into model improvement, rather than treating human correction as an exception.

Another priority is the development of open benchmarks and realistic testbeds. Many published results rely on simulation or limited laboratory data, which makes comparison difficult and deployment risky. Shared datasets, digital twins, and emulation environments can help researchers evaluate models under realistic fault conditions, including gradual degradation, intermittent errors, and multi-vendor interoperability. Policy support for such infrastructure can accelerate innovation while protecting commercial confidentiality. Finally, research should examine the environmental and equity consequences of automated fault detection. Metrics such as avoided outages, reduced truck rolls, energy savings, and improved service availability in underserved regions should be included in evaluation frameworks. This would align technical progress with public-interest outcomes.

10. Conclusion

Machine learning-enabled fault detection in high-speed fiber-optic networks offers significant potential to improve reliability, reduce downtime, and support increasingly demanding digital services. However, the decisive challenges are not limited to algorithmic accuracy. They include observability, data governance, latency constraints, interpretability, cost, vendor interoperability, workforce capacity, sustainability, and equitable access to resilient infrastructure. Optical networks are complex socio-technical systems in which faults emerge from physical, logical, and organizational interactions. Machine learning can help operators see patterns that traditional alarms miss, but it must be embedded in architectures and institutions that make uncertainty visible and accountability traceable.

The most promising path forward is not the replacement of human expertise but the augmentation of it. Hybrid diagnostic platforms that combine physics-based reasoning, data-driven inference, and operator feedback are more likely to succeed than fully autonomous black boxes. Governance frameworks should specify automation levels, model monitoring, rollback procedures, and transparency requirements according to the criticality of the service. Policy should encourage open standards, shared testbeds, and inclusive deployment so that the benefits of advanced fault detection do not accrue only to the most privileged networks. If these conditions are met, machine learning can contribute not only to faster fault detection but also to a more robust, sustainable, and fair optical infrastructure.

References

1. Agrawal, G. P. (2021). Fiber-optic communication systems (5th ed.). Wiley.
2. Cisco. (2020). Cisco annual internet report (2018–2023) white paper. Cisco.
3. ITU-T. (2020). G.694.1: Spectral grids for WDM applications: DWDM frequency grid. International Telecommunication Union.
4. Essiambre, R.-J., Kramer, G., Winzer, P. J., Foschini, G. J., & Goebel, B. (2010). Capacity limits of optical fiber networks. *Journal of Lightwave Technology*, 28(4), 662–701.
5. Winzer, P. J., & Essiambre, R.-J. (2006). Advanced optical modulation formats. *Proceedings of the IEEE*, 94(5), 952–985.
6. Poggiolini, P. (2012). The GN model of non-linear propagation in uncompensated coherent optical systems. *Journal of Lightwave Technology*, 30(24), 3857–3879.
7. Kilper, D. C., Atkinson, G., Korotky, S. K., Goyal, P., Vetter, P., Suvakovic, D., & Blouza, A. (2011). Power trends in communication networks. *IEEE Journal of Selected Topics in Quantum Electronics*, 17(2), 275–284.
8. Vasseur, J. P., & Pickavet, M. (2008). Network recovery: Protection and restoration of optical, SONET-SDH, IP, and MPLS. Morgan Kaufmann.
9. Mukherjee, B. (2006). Optical WDM networks. Springer.
10. Fawaz, H. I., Forestier, G., Weber, J., Idoumghar, L., & Muller, P.-A. (2019). Deep learning for time series classification: A review. *Data Mining and Knowledge Discovery*, 33(4), 917–963.
11. Musumeci, F., Rottondi, C., Nag, A., Macaluso, I., Zibar, D., Ruffini, M., & Tornatore, M. (2019). An overview on application of machine learning techniques in optical networks. *IEEE Communications Surveys & Tutorials*, 21(2), 1383–1408.
12. Mata, J., de Miguel, I., Durán, R. J., Merayo, N., Singh, S. K., Jukan, A., & Chamanian, M. (2018). Artificial intelligence (AI) methods in optical networks: A comprehensive survey. *Optical Switching and Networking*, 28, 43–57.
13. Saif, W. S., Esmail, M. A., Ragheb, A. M., Alshawi, T. A., & Alshebeili, S. A. (2021). Machine learning techniques for optical performance monitoring and modulation format identification: A survey. *IEEE Communications Surveys & Tutorials*, 23(4), 2272–2304.
14. Rottondi, C., Barletta, L., Giusti, A., & Tornatore, M. (2018). Machine-learning method for quality of transmission prediction of unestablished lightpaths. *Journal of Optical Communications and Networking*, 10(2), A286–A297.
15. Zhou, Z., Wei, J., Luo, Y., Clark, K. A., Sillekens, E., Deakin, C., ... & Liu, Z. (2023). Communications with guaranteed bandwidth and low latency using frequency-referenced multiplexing. *Nature Electronics*, 6(9), 694–702.
16. Zibar, D., Piels, M., Jones, R., & Schäffer, C. G. (2016). Machine learning concepts in optical communication. *Journal of Lightwave Technology*, 34(6), 1446–1457.
17. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
18. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.

19. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215.
20. Winner, L. (1980). Do artifacts have politics? *Daedalus*, 109(1), 121–136.