

Trade-Offs of AI-based Optical Performance Monitoring and Sensing Techniques

Downloaded from: https://research.chalmers.se, 2026-11-22 18:11 UTC

Citation for the original published paper (version of record):

Natalino Da Silva, C., Kaeval, K., Wosinska, L. (2026). Trade-Offs of AI-based Optical Performance Monitoring and Sensing Techniques. 52nd European Conference on Optical Communications, ECOC 2026

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

research.chalmers.se offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all kind of research output: articles, dissertations, conference papers, reports etc. since 2004. research.chalmers.se is administrated and maintained by Chalmers Library

(article starts on next page)

Trade-Offs of AI-based Optical Performance Monitoring and Sensing Techniques

Carlos Natalino[(1)], Kaida Kaeval[(2)], Lena Wosinska[(1)]

[(1)] Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden, wosinska@chalmers.se
[(2)] Tallinn University of Technology, Tallinn, Estonia.

Abstract *Modern optical networks leverage AI/ML for performance monitoring and fiber sensing across Φ-OTDR/DAS and SOP technologies. We present a unified data-rate framework, showing that monitoring precision choices impact the generated data volumes over nine orders of magnitude, and analyze AI/ML processing trade-offs across these configurations.* *© 2026 The Author(s)*

Introduction

Modern optical networks target both high availability and low operational margins, making proactive monitoring essential for early anomaly detection and failure prevention [1]. Meanwhile, integrated sensing and communication (ISAC) enables telecom fibers to serve as distributed environmental sensors capable of detecting vibrations, strain, etc. along the fiber [2, 3]. Artificial intelligence (AI) and machine learning (ML) techniques have become essential for the automation of failure management, from soft-failure detection to root-cause identification and predictive maintenance [4, 5].

However, its effectiveness depends on data quality, availability, and representativeness. While accuracy and precision are set by the sensing hardware, representativeness and availability require continuous collection at short intervals. Transmission and storage of such high-capacity data is challenging, forcing trade-offs between accuracy, throughput, and latency [6]. A systematic understanding of these trade-offs is key to dimensioning the monitoring pipeline efficiently.

This paper presents a framework for data-rate evaluation across optical performance monitoring (OPM) and sensing devices. We also quantify the trade-offs between monitoring precision and data volume, monitoring frequency and localization accuracy, and AI/ML model complexity and inference time. Together, they provide a principled basis for estimating realistic telemetry and compute budgets for AI-based monitoring pipelines.

Unified Data-Rate Framework

Network telemetry has shifted away from poll-based protocols such as simple network management protocol (SNMP), which can introduce minutes of latency between a physical state change and its observation. Modern deployments instead rely on push-based streaming protocols built on gRPC remote procedure call (gRPC), such as gRPC network management interface (gNMI) [7].

At the same time, optical sensing use cases have surged in relevance, ranging from fiber integrity monitoring to environmental sensing, generating large volumes of data that must be reliably collected and delivered to processing pipelines, in addition to supporting multi-domain services such as optical-spectrum-as-a-service (OSaaS) [8]. The telemetry-as-a-service (TaaS) paradigm further decouples stateful collectors from stateless injectors, allowing seamless integration of intermediate processing steps [9].

Fig. 1 shows the three-layer monitoring architecture considered in this work. At the bottom, the optical network physical infrastructure is instrumented with four representative data sources: optical performance monitoring (OPM) extracted from the data channel, state of polarization (SOP) monitoring, conventional optical time-domain reflectometry (OTDR), and distributed acoustic sensing based on phase-sensitive OTDR (Φ-OTDR). For clarity, each modality is drawn on a distinct link/path. However, in deployed networks these channels can also co-propagate on the same fibers with data traffic, and multiple modalities can coexist on a single link/path. The resulting raw monitoring

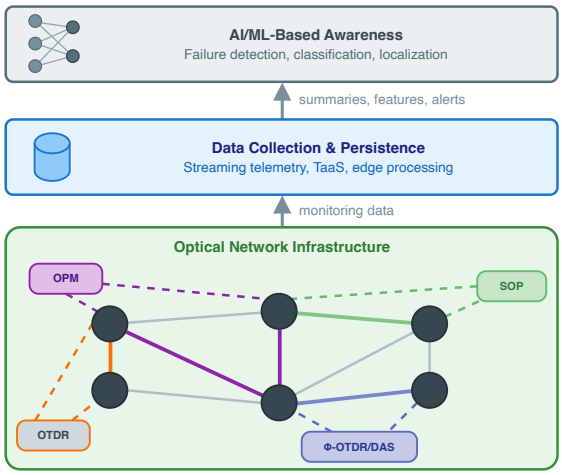


Fig. 1: General monitoring architecture. Required data rates from the physical layer span nine orders of magnitude depending on the technology and configuration.

data moves up to a data collection and persistence layer built on streaming telemetry, TaaS, etc. Generated summaries, features, and alerts are forwarded to the AI/ML-based awareness layer responsible for the awareness tasks. The awareness layer can run centralized, distributed, or hybrid across edge and data-center sites, either standalone or embedded in a network controller.

The data rates produced by the telemetry and sensing devices vary by several orders of magnitude depending on the sensing modality and its configuration. In the following, we unify the data-rate analysis across technologies through a common framework: the total required data rate is determined by the number of sensing channels, the sampling frequency, and a multiplier for the detection scheme. These parameters define the achievable spatial resolution. Throughout, c [$\mathrm{m/s}$] denotes the speed of light in vacuum, n [-] the fiber refractive index, L [m] the fiber length, and f_s [Hz] the sampling frequency.

OTDR-based methods. For backscattering-based methods (OTDR, Φ-OTDR, distributed acoustic sensing (DAS)), the maximum pulse repetition rate (PRR) is constrained by the fiber round-trip time:

$$\mathrm{PRR}_{\max} = \frac{c}{2nL} \tag{1}$$

where $\mathrm{PRR}_{\max}$ [Hz] is the maximum pulse repetition rate [10]. For a $50\,\mathrm{km}$ fiber, $\mathrm{PRR}_{\max} \approx 2\,\mathrm{kHz}$, limiting the acoustic sensing bandwidth to $1\,\mathrm{kHz}$ by the Nyquist theorem.

The data rate for Φ-OTDR/DAS is:

$$R_\Phi = \frac{L}{\Delta x} \cdot \mathrm{PRR} \cdot K \tag{2}$$

where R_Φ [Sa/s] is the raw sample rate, Δx [m] is the spatial resolution, and K [-] accounts for the detection scheme ($K{=}1$ for direct detection, $K{=}2$ for coherent I/Q) [10].

A heterodyne Φ-OTDR interrogator digitizing the beat signal at $250\,\mathrm{MSa/s}$ with 16-bit resolution produces approximately $500\,\mathrm{MB/s}$ [11]. Standard direct-detect OTDR instead operates in batch mode, averaging many pulses per on-demand scan, falling short if real-time fiber certification and hard-fault localization are required, [1, 4] but places its time-averaged data rate orders of magnitude below Φ-OTDR/DAS.

SOP monitoring. The per-channel data rate is set by the sampling frequency and the number of polarization parameters recorded per sample:

$$R_{\mathrm{SOP}} = f_s \cdot N_S \tag{3}$$

where R_{SOP} [val/s] is the per-channel data rate and N_S [-] is the number of polarization parameters per sample, with $N_S{=}3$ for the normalized Stokes parameters (s_1, s_2, s_3) or $N_S{=}4$ when the total power S_0 is also recorded. Unlike the OTDR-family modalities discussed above, SOP monitoring does not require a dedicated interrogator: the Stokes vector can be extracted directly from the digital signal processing pipeline of a coherent transceiver, utilizing infrastructure that is already deployed on every lit channel [2, 12]. As a consequence, monitoring precision is defined by the choice of sampling rate that can span many orders of magnitude. Field studies of metropolitan hazards such as earthquakes, mechanical vibrations from construction activity or excavators, and cable flaps caused by maintenance or environmental stress operate at average sampling rates on the order of $17\,\mathrm{Hz}$ [2], and a few tens of Hz are sufficient to capture the spectral signature of these events [12]. Alternatively, symbol-rate monitoring at 64 Gbaud can generate up to $256\,\mathrm{Gval/s}$ per channel.

Because SOP is accumulated over the entire fiber, a single-ended deployment cannot localize events along the link [2, 12, 13]. For spatial resolution, bidirectional SOP monitoring partially recovers spatial information, with a localization precision bounded by:

$$\Delta x_{\mathrm{SOP}} = \frac{c}{2n \cdot f_s} \tag{4}$$

where Δx_{SOP} [m] is the achievable spatial precision and f_s is the per-end sampling rate, under the assumption of perfect bidirectional symmetry in fiber length and component delay.

Trade-off Analysis

Spatial Precision vs. Data Volume. The total data rate of the collected monitoring data is driven by three precision dimensions: spatial resolution (finer spacing increases the number of channels linearly), temporal resolution (higher sampling rates capture faster phenomena), and amplitude precision (bit depth and detection scheme). These three dimensions multiply to yield the total data rate, forming a "precision cube" in which any improvement along one axis proportionally increases data volume.

Fig. 2 illustrates this across specific configurations. Data rates span from $51\,\mathrm{val/s}$ for sub-sampled SOP telemetry to $256\,\mathrm{Gval/s}$ for symbol-rate SOP, with DAS and Φ-OTDR occupying intermediate positions. An inset marks the $\approx 1\,\mathrm{Gval/s}$ ceiling of a PCIe 5.0 x4 interface (commonly adopted by enterprise-grade storage devices) at 16-bit per value. Only symbol-rate SOP overshoots this budget, by more than two orders of magnitude, requiring dedicated interconnect or edge-side compression.

Fig. 3 reveals the same inverse inverse relationship across all technologies. Bidirectional SOP

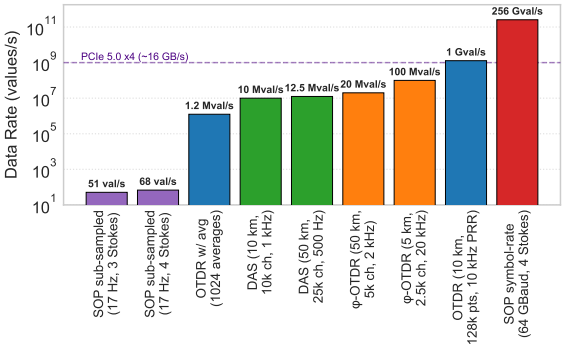


Fig. 2: Data rate comparison across usual and futuristic sensing configurations.

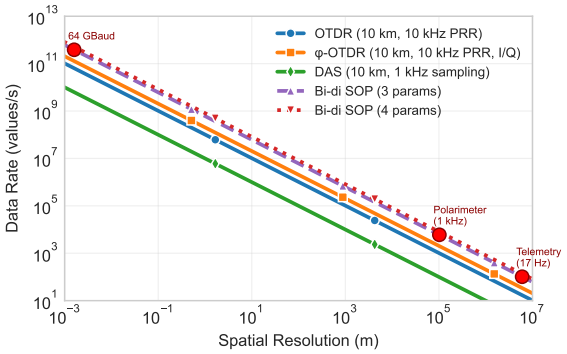


Fig. 3: Spatial resolution vs. data rate across sensing technologies. Bidirectional SOP lines show both 3-Stokes and 4-Stokes configurations. Key operating points are marked.

at $17\,\mathrm{Hz}$ telemetry yields $\sim 6000\,\mathrm{km}$ precision, providing no effective localization at metropolitan distances, while symbol-rate monitoring at $64\,\mathrm{Gbaud}$ reaches sub-millimeter resolution at the cost of $256\,\mathrm{Gval/s}$ per channel. Φ-OTDR provides meter-scale resolution natively, but at the cost of dedicated interrogator hardware, higher data volumes, and specific pulse-length configuration.

The application determines which precision dimension to prioritize: seismic monitoring requires temporal bandwidth above $100\,\mathrm{Hz}$ but tolerates $10\,\mathrm{m}$ spatial resolution [2], while perimeter awareness demands finer spatial localization [14]. SOP monitoring offers a cost-effective alternative when fine-grained localization is not required, since the Stokes vector can be read directly from the digital signal processing (DSP) pipeline of coherent transceivers. Early anomaly detection over a live metropolitan link is feasible at telemetry-level SOP sampling rates available in commercial transceivers [2], and accurate classification of harmful fiber events using only SOP traces collected on an outdoor testbed are possible [15].

Monitoring Frequency vs. Temporal Localization. The temporal precision with which an AI model can pinpoint the onset of a soft failure is bounded by the nature of the failure development time and by how often the underlying monitoring data is collected. Failures that occur between two consecutive samples remain unobserved, and the predicted time of onset cannot be resolved finer than the probe period itself. Hence, tighter temporal localization demands faster probing, which in turn multiplies the rate at which telemetry must be collected, transported, and processed.

AI/ML Model Complexity vs. Inference Time. The inference time of the chosen model imposes two hard constraints on the monitoring pipeline. First, it caps the rate at which a single inference pipeline can consume incoming monitoring data: if the per-sample inference time exceeds the interval between samples, the pipeline cannot keep up, and sustaining the full data rate requires running several inference instances in parallel at a proportional compute cost. Second, inference time adds directly to the delay between the physical event and the moment a detection becomes actionable, which matters for time-sensitive tasks such as intrusion alerting or incident response. Lightweight classifiers deliver sub-millisecond inference and can sustain single-pipeline processing even on high-rate modalities such as Φ-OTDR/DAS [14, 16]. Heavier deep architectures such as Vision Transformers reach higher classification accuracy at the cost of orders of magnitude more compute per inference [17], forcing either parallel deployment or a coarser effective sampling of the incoming data. Foundation models applied in a zero-shot regime remove the need for labeled training data but incur inference latencies on the order of seconds [18], restricting them to slow-timescale tasks. As sensing data rates grow, only lightweight models remain feasible.

Conclusions

We presented a unified data-rate framework for optical performance monitoring and sensing across OTDR, Φ-OTDR/DAS, and SOP modalities, and used it to expose three interrelated trade-offs. Monitoring precision decomposes into spatial, temporal, and amplitude axes whose contributions multiply rather than add, so any gain along one axis inflates the data volume proportionally. Temporal fault localization is bounded by the probing frequency, which itself drives the data rate that collection and processing must sustain. AI model choice is constrained by inference time through both a throughput ceiling and an event-to-detection latency floor, and both constraints tighten as sensing data rates grow. The optimal monitoring configuration is therefore application-driven, trading precision and responsiveness against the compute and telemetry budget available along the pipeline. Future work should address the co-design of sensing hardware, telemetry transport, AI inference, and OSaaS [8], together with scalable telemetry architectures such as TaaS [9] to support these workflows.

Acknowledgements

This work was funded by the European Commission's Horizon Europe research and innovation programme within the SNS JU project ECO-eNET (grant agreement no. 101192006) and the Digital Europe Programme (101127973) through the 5G-TACTIC project.

References

[1] F. Musumeci and M. Tornatore, "Failure management in optical networks with ML: A tutorial on applications, challenges, and pitfalls [Invited]", *Journal of Optical Communications and Networking*, vol. 17, no. 8, pp. C144–C155, Aug. 2025. DOI: 10.1364/JOCN.551910.

[2] C. J. Carver and X. Zhou, "Polarization sensing of network health and seismic activity over a live terrestrial fiber-optic cable", *Communications Engineering*, vol. 3, no. 1, p. 91, Jul. 2024. DOI: 10.1038/s44172-024-00237-w.

[3] A. Rode et al., "Machine learning opportunities for integrated polarization sensing and communication in optical fibers", *Optical Fiber Technology*, vol. 90, p. 104 047, Mar. 2025. DOI: 10.1016/j.yofte.2024.104047.

[4] D. Wang, C. Zhang, W. Chen, H. Yang, M. Zhang, and A. P. T. Lau, "A review of machine learning-based failure management in optical networks", *Science China Information Sciences*, vol. 65, no. 11, p. 211 302, Oct. 2022, ISSN: 1869-1919. DOI: 10.1007/s11432-022-3557-9. Accessed: Jan. 14, 2026.

[5] S. Cruzes, "Failure Management Overview in Optical Networks", *IEEE Access*, vol. 12, pp. 169 170–169 193, 2024. DOI: 10.1109/ACCESS.2024.3498704.

[6] S. Cruzes, "Telemetry and Agentic AI: Foundations for Optical Network Automation", *IEEE Access*, vol. 14, pp. 8800–8838, 2026. DOI: 10.1109/ACCESS.2025.3649768.

[7] P. Castoldi et al., "Shaping the future of optical networks by integrating SDN, telemetry, and AI [Invited]", *Journal of Optical Communications and Networking*, vol. 17, no. 7, pp. C51–C61, Jul. 2025. DOI: 10.1364/JOCN.553843. Accessed: Jan. 22, 2026.

[8] K. Kaeval, K. Grobe, and J.-P. Elbers, "Operation of optical spectrum as a service in disaggregated and multi-operator environments [invited]", *Journal of Optical Communications and Networking*, vol. 17, no. 1, A46–A58, Jan. 2025. DOI: 10.1364/JOCN.534118.

[9] C. Natalino, K. Kaeval, H. J. Kerm, T. Järvelill, M. Jasper, and P. Monti, "Telemetry-as-a-Service: Decoupling collection from injection for scalable network monitoring", in *2026 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, Jun. 2026.

[10] M. R. Fernández-Ruiz, L. Costa, and H. F. Martins, "Distributed Acoustic Sensing Using Chirped-Pulse Phase-Sensitive OTDR Technology", *Sensors*, vol. 19, no. 20, Oct. 2019. DOI: 10.3390/s19204368.

[11] F. Yu et al., "Data Reduction in Phase-Sensitive OTDR with Ultra-Low Sampling Resolution and Undersampling Techniques", *Sensors*, vol. 22, no. 17, Aug. 2022. DOI: 10.3390/s22176386.

[12] S. Pellegrini et al., "Overview on the state of polarization sensing: Application scenarios and anomaly detection algorithms", *Journal of Optical Communications and Networking*, vol. 17, no. 2, A196–A209, Feb. 2025. DOI: 10.1364/JOCN.537881.

[13] K. Abdelli, M. Lonardi, J. Gripp, S. Olsson, F. Boitier, and P. Layec, "Risky event classification leveraging transfer learning for very limited datasets in optical networks", *Journal of Optical Communications and Networking*, vol. 16, no. 7, pp. C51–C68, Jul. 2024. DOI: 10.1364/JOCN.517529.

[14] A. Tomasov et al., "Enhancing Perimeter Protection using Φ-OTDR and CNN for Event Classification", in *28th International Conference on Optical Fiber Sensors (2023)*, Nov. 2023, W4.39. DOI: 10.1364/OFS.2023.W4.39.

[15] L. Sadighi, S. Karlsson, C. Natalino, L. Wosinska, M. Ruffini, and M. Furdek, "Deep Learning for Detection of Harmful Events in Real-World, Noisy Optical Fiber Deployments", *Journal of Lightwave Technology*, vol. 43, no. 13, pp. 6092–6101, Jul. 2025. DOI: 10.1109/JLT.2025.3557748.

[16] G. Malik et al., "Resilient Anomaly Detection in Fiber-Optic Networks: A Machine Learning Framework for Multi-Threat Identification Using State-of-Polarization Monitoring", *AI*, vol. 6, no. 7, p. 131, Jul. 2025. DOI: 10.3390/ai6070131.

[17] K. Abdelli et al., "Vision Transformers for Anomaly Classification and Localization in Optical Networks Using SOP Spectrograms", *Journal of Lightwave Technology*, vol. 43, no. 4, pp. 1902–1914, Feb. 2025. DOI: 10.1109/JLT.2024.3519755.

[18] K. Abdelli, "Zero-shot forecasting of optical network telemetry using large language models", *Journal of Optical Communications and Networking*, vol. 17, no. 12, pp. 1148–1160, Dec. 2025. DOI: 10.1364/JOCN.575791.