Optical Fiber Technology 101 (2026) 104715

Contents lists available at ScienceDirect

Optical Fiber Technology

journal homepage: www.elsevier.com/locate/yofte

Regular Articles

Deep-learning-based tool for anomaly detection and optical signal-to-noise ratio/Channel Spacing estimation in Elastic Optical Networks

Kevin David Martinez-Zapata *, Jhon James Granada-Torres

Universidad de Antioquia, Facultad de Ingeniería. GITA lab, calle 67 No 52-108, Medellín, 050004, Antioquia, Colombia

ARTICLE INFO

Keywords:
Anomaly detection
Parameter estimation
Elastic optical networks
Machine learning
Inter-channel interference

ABSTRACT

Detecting anomalies and estimating system parameters are crucial for maintaining service quality in gridless Elastic Optical Networks (EONs). The absence of fixed Channel Spacing (ChSp) and the need for flexible resource allocation make anomaly detection and parameter estimation particularly challenging, especially when ChSp is close to or smaller than the Channel Bandwidth (ChBw). In this work, we propose a deep learning (DL)-based optical performance monitoring tool built upon a common neural architecture that can be parametrically configured for either classification or regression, depending on the target monitoring task. When configured in classification mode, the model detects anomalies, such as excessive noise levels and spectral overlap, whereas in regression mode, it provides an accurate estimate of ChSp or OSNR. This task-oriented configurability allows the same architectural framework to be applied across multiple monitoring functions without requiring fundamentally different model designs. We experimentally validated our method in 3 × 16- and 3 × 32-GBd 16-Quadrature-Amplitude-Modulation (QAM) Nyquist Wavelength-Division Multiplexing (WDM) systems. Our tool achieved approximately 98% and 89% accuracy in binary spectral-overlap classification for 16- and 32-GBd cases, respectively. ChSp estimation errors were limited to 0.3 GHz (16-GBd) and 0.9 GHz (32-GBd). Detection of low- and high-OSNR levels achieved approximately 96% and 81%–85% accuracy for 16-GBd and all the 32-GBd scenarios, respectively. Estimation errors less than 0.5 dB for the 16-GBd scenario and 0.8–1.6 dB for all the 32-GBd scenarios were achieved. These results demonstrate the robustness and versatility of the proposed OPM framework, which operates exclusively at the symbol level, thereby avoiding access to higher-rate signal samples or modifications to the Digital Signal Processing (DSP) chain of conventional coherent receivers, making it well-suited for future dynamic gridless optical transmission systems.

1. Introduction

The exponential growth in data traffic and the increasing demand for high-speed connectivity have driven the evolution of Elastic Optical Networks (EONs), where dynamic resource allocation and gridless spectrum management are crucial for meeting future bandwidth requirements [1]. The spectral flexibility in EONs based on wavelength-division multiplexing (WDM) is enabled by reconfigurable optical add-drop multiplexers (ROADMs) that dynamically allocate optical channels [2]. However, the dynamic and gridless nature of EONs introduces significant challenges for optical performance monitoring (OPM), particularly regarding channel spacing (ChSp) and optical signal-to-noise ratio (OSNR) [3,4]. On one hand, future gridless networks must estimate the spectral separation between adjacent channels at transit nodes to determine optimal channel placement and the required spectral width for the “adding” and “dropping” processes. This presents a challenge, as each channel may occupy a different channel bandwidth (ChBw) depending on its symbol rate and roll-off factor. On the other hand, in gridless scenarios with reduced or absent guard bands, where ChSp is close to or even lower than the ChBw, estimating OSNR becomes particularly challenging due to the difficulty of isolating the noise floor from the signal [3,5,6].

While EONs greatly enhance spectral efficiency through flexible channel allocation, their performance is fundamentally constrained by two interrelated physical-layer impairments: inter-channel interference (ICI) and OSNR degradation. Previous experimental studies have shown that linear ICI can severely degrade system performance, with Bit Error Rate (BER) strongly dependent on both OSNR and ChSp, including scenarios with spectral overlap [7]. These results highlight that ICI is not merely a secondary impairment, but a dominant factor that must be accurately quantified under gridless operation. Simultaneously, the reduction in ChSp diminishes the effective OSNR margin, as Amplified Spontaneous Emission (ASE) noise accumulates across adjacent

* Corresponding author.
E-mail address: kevin.martinez1@udea.edu.co (K.D. Martinez-Zapata).

https://doi.org/10.1016/j.yofte.2026.104715
Received 2 February 2026; Received in revised form 19 April 2026; Accepted 10 June 2026

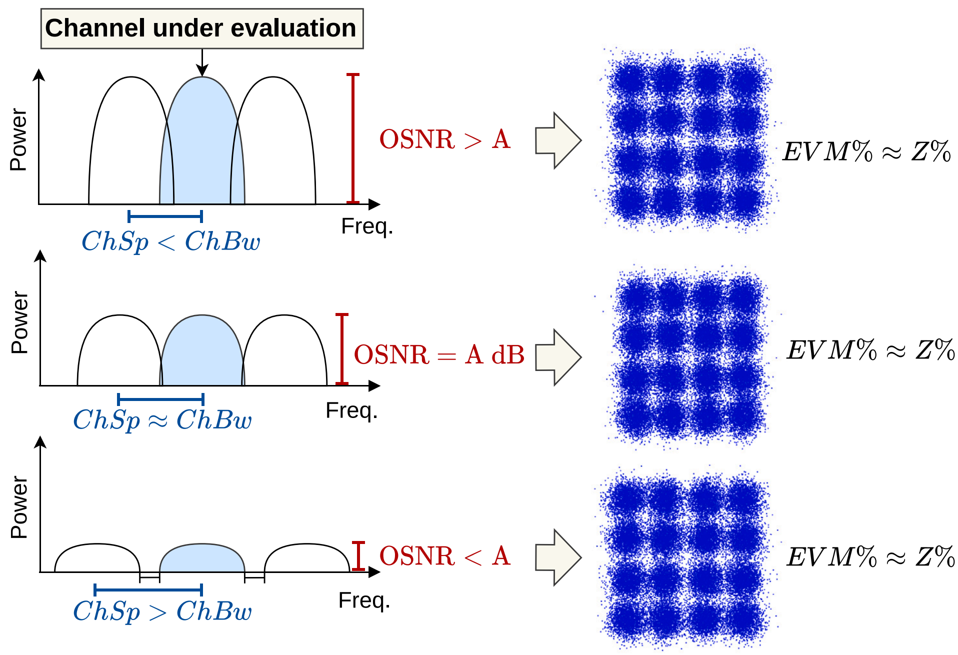

Fig. 1. Spectrum diagrams of three configurations, each depicting three optical channels with different spectral separations and OSNR values. Some channels are affected by linear ICI, while others remain unaffected. The IQ constellation diagrams indicate that all three configurations yield similar EVM values.

channels [8]. Traditional monitoring systems struggle to isolate these effects; OSNR measurements (e.g., via spectral analysis) often misinterpret linear ICI as additional noise, while ICI detection methods typically rely on knowledge of neighboring channel parameters [9]. This interdependence results in a diagnostic blind spot within networks, where neither impairment can be accurately assessed without prior knowledge of the other.

For instance, as illustrated in Fig. 1, the received symbols in the In-Phase/Quadrature (IQ) constellation diagrams exhibit significant visual similarity across different combinations of ChSp and OSNR levels. For example, a configuration with guard bands (no spectral overlap) and low OSNR can produce a constellation diagram nearly indistinguishable from one with narrow spectral overlap and high OSNR, as both scenarios result in comparable % Error Vector Magnitude (% EVM) values. This ambiguity poses a critical challenge for traditional OPM methods, as it complicates the accurate discrimination between noise-dominated and interference-dominated impairments. In addition to linear ICI and ASE noise, fiber nonlinearities, such as those induced by the Kerr effect, may further distort the received signal, particularly at higher launch powers. These effects introduce additional phase and amplitude perturbations, increasing the complexity of performance monitoring in practical systems. Also, given the considerable computational cost of temporal waveform analysis and Fourier transformations, this underscores the necessity of OPM at the symbol level. Consequently, conventional approaches based on optical spectrum analyzers (OSAs) or simple % EVM metrics fail to provide reliable OSNR or ChSp estimates in dynamic gridless networks.

Recent studies [1–4,10,11] have used machine learning (ML) techniques, including deep learning (DL) architectures, to estimate OSNR at specific ChSp conditions [1,2,10]. Furthermore, in [3,11], the authors introduced DL-based methods to estimate ChSp in gridless networks, although at a fixed OSNR value.

In this context, this work proposes a configurable DL-based optical performance monitoring methodology built upon a common feature extraction and processing pipeline. By leveraging 2D histogram representations of IQ constellations, the approach enables the discrimination between linear ICI and ASE noise, even under spectral overlap conditions. The same methodological framework is systematically applied to different monitoring tasks, including anomaly detection and the estimation of OSNR and ChSp, across different symbol rates and operating regimes. As an extension of our preliminary conference work [12], which focused exclusively on OSNR estimation under constrained scenarios, this work expands the methodology to additional monitoring tasks, including ChSp estimation and anomaly detection, along with a more comprehensive experimental validation. In contrast to existing approaches, which are typically limited to single-task estimation under fixed conditions and often neglect spectral overlap effects, the proposed framework is evaluated under more realistic gridless scenarios, including sub-Nyquist channel spacing and varying system configurations. This enables a more consistent and robust assessment of optical performance monitoring while maintaining a symbol-level implementation that avoids modifications to the DSP chain. A detailed comparison with representative methods from the literature is provided in Table 1, highlighting differences in monitoring scope, feature representation, dataset diversity, and performance metrics, as well as the corresponding advantages and limitations of each approach.

The remainder of this article is organized as follows: Section 2 describes the experimental setup and the data generation process. Subsequently, Section 3 describes the development of the monitoring tool, from data preprocessing to the DL strategies that enable the estimation and detection of anomalies in two gridless Nyquist-WDM systems with 16- and 32-GBd 16-QAM. Section 4 presents the results of the DL-based monitoring tool along with a discussion on deployment feasibility. Finally, Section 5 outlines the conclusions of this work and discusses future research directions.

2. Experimental setup

The dataset used to demonstrate the functionality of the proposed tool was derived from two experimental setups, ES1 and ES2, as detailed in Ref. [7]. Both setups employ a gridless Nyquist-WDM system using 16-QAM modulation with baud rates of 16-GBd for ES1 and 32-GBd for ES2. Each setup incorporates three tunable laser configurations, with linewidths of 100 kHz for ES1 and 25 kHz for ES2, enabling adjustment of the spectral separation between the channels. Laser modulation is achieved through dual-polarization IQ modulators, which are

Table 1
Architectural comparison of representative ML-based approaches for OPM.

Ref.	ML Architecture & core function	Feature representation	Dataset size & transmission Profile	Performance levels	Pros	Cons
Thrane et al. (2017) [4]	SVM (Classification) & Shallow NN (Regression) for Modulation Format Identification (MFI) & OSNR	1D Features (Mean/Variance extracted from Eye Diagrams)	Tested on specific 32-GBd QAM. Fixed channel spacing assumed. OSNR: 4–30 dB	MFI Acc: 94%. OSNR Error: 0.2 dB (only within a specific 4–17 dB range)	Fast execution due to very simple, low-dimensional features. Excellent estimation within constrained OSNR bands	Discards all phase info. Cannot distinguish between CD, PMD, and severe linear ICI. Fails catastrophically at OSNR extremes due to feature blending (1.2 dB)
Wan et al. (2019) [2]	Multi-Task Learning ANN (MTL-ANN) for joint OSNR & MFI	1D Amplitude Histograms (AH)	Evaluated on discrete OSNR ranges. Fixed grid, no overlapping channel characterization	OSNR MSE: 0.12 dB. MFI Acc: 100% (on simple formats)	Reduces computation time by >40% via shared hidden layers for multiple tasks	1D Histograms obliterate phase-quadrature covariance, failing completely to capture the spatial asymmetrical deformations of linear ICI
Locatelli et al. (2021) [11]	Curve Fitting + ML Regression for filter bandwidth & central frequency shift estimation (ingress node) and ASE noise estimation (egress node)	Optical spectrum magnitude from low-cost Optical Channel Monitors (OCMs)	Sim. (16 cases) + Exp. (7 cases); DWDM/flex-grid; filter 6 dB BW: 36.5–38.5 GHz; freq. shift: ±2 GHz; VOA-emulated ASE noise	Filter freq. shift MAE: <0.63 GHz (sim.)/<1 GHz (exp.); filter 6 dB BW MAE: <1 GHz (exp.); ASE noise MSE: 0.014 (exp.)	Dual ingress/egress monitoring with low-cost OCMs. Achieves 84% and 75% relative accuracy gain for freq. shift and BW vs. datasheet values. Applicable to disaggregated networks	Spectral-domain input only - no symbol-level IQ analysis. No MFI or direct OSNR waveform estimation. Egress ASE relative error remains high (78%–91%). Limited filter parameter space
Saif et al. (2022) [1]	SVR for OSNR & Cromatic Dispersion (CD) estimation	2D-IQH (compressed via 2D-DFT/DCT transforms)	Exp. 7 × 20-GBd DP-QPS system, with OSNR in the range of 12–24, with a step of 2 dB	OSNR R^2> 0.98. CD R^2> 0.97	Superior dimensionality reduction reduces computational complexity at the receiver	Emulated chromatic dispersion. Experimentally validated only for QPSK. No extreme negative guard band (overlap) analysis. Only 20-GBd
Hraghi et al. (2023) [3]	Feed-Forward Neural Network (FFNN) for channel spacing estimation	1D projections of IQ histograms onto diagonal/horizontal axes	Limited OSNR ranges (17 to 19 dB in the experimental setup), focused exclusively on spacing. 3 × 52-GBd 16-QAM	Spacing RMSE: 0.32 GHz (sim.)/1.26 GHz (exp.)	Specifically targets the critical issue of linear inter-channel crosstalk in dense WDM	Computationally isolated (does not estimate OSNR). Experimental error is relatively high (1.26 GHz) and lacks robust validation under severe overlap. Limited OSNR range
Our Framework	Configurable CNN, FCNN, AE-FCNN for spectral overlap and low-OSNR detection (anomaly detection), and system parameters estimation (ChSp and OSNR)	2D Histograms of IQ Symbols (Gaussian filtered & Normalized).	Exhaustive gridless sweep (>80 scenarios, >100k symbols each). Spacing: 15–18 GHz, including single channel (16GBd) and 28.5–50 GHz (32GBd). Includes severe spectral overlap (ChSp < ChBw). OSNR: 14–40 dB	Spacing RMSE: <0.6 GHz. Overlap Acc: ≈98%. OSNR MAE: <2 dB (blind). R2>0.8 for all scenarios	Unified pipeline seamlessly swaps between **classification and regression**. Good blind estimation capability. Validated on **negative guard bands**. Characterization across a **wide range of ChSp and OSNR**. It works **at the symbol level**	The method was experimentally validated using two spectrally adjacent, equally spaced 16QAM-modulated channels

driven by four digital-to-analog converters (DACs). The DACs operate at sampling rates of 64 GSa/s for ES1 and 88 GSa/s for ES2, with electrical bandwidths of 14 GHz and 32 GHz, respectively. Both center and side frequency channels are generated by different DACs to produce uncorrelated signals.

A roll-off factor of 0.1 is used to implement Nyquist pulse shaping, resulting in ChBw of 17.6 GHz for the experimental setup ES1 and 35.2 GHz for ES2. The dataset per setup includes over 80 scenarios, each containing more than 40k received symbols. The variable ChSp is incremented by 0.5 GHz, ranging from 15 to 18 GHz for ES1 and 30 to 36 GHz for ES2. This includes a ChSp of 18 GHz (in a non-overlapping scenario) and single-channel cases for ES1, as well as a ChSp of 37.5 GHz (in a non-overlapping scenario) for ES2.

The experiments evaluate the impact of ICI and include both back-to-back (B2B) and optical fiber transmission configurations, as illustrated in Fig. 2. ASE noise is introduced before the receiver by saturating an optical amplifier, resulting in an OSNR from 14 to 40 dB. Transmission over 270 km of standard single-mode fiber is included for ES2. Erbium-doped fiber amplifiers (EDFAs) were placed after 90 km links. The launch powers for the transmission were set to 0 dBm

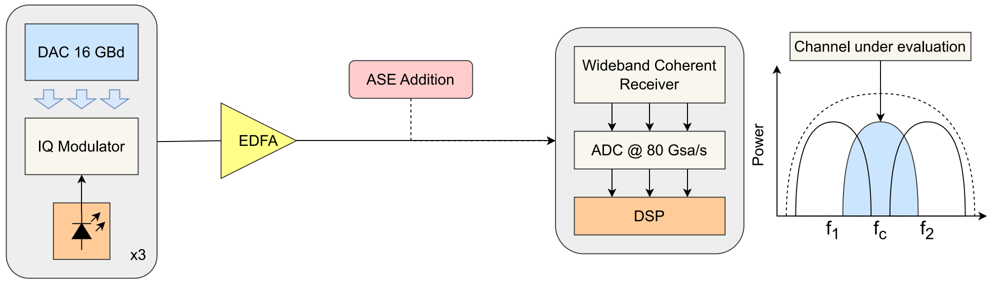

(a) 3 × 16-GBd 16-QAM Nyquist-WDM system (ES1).

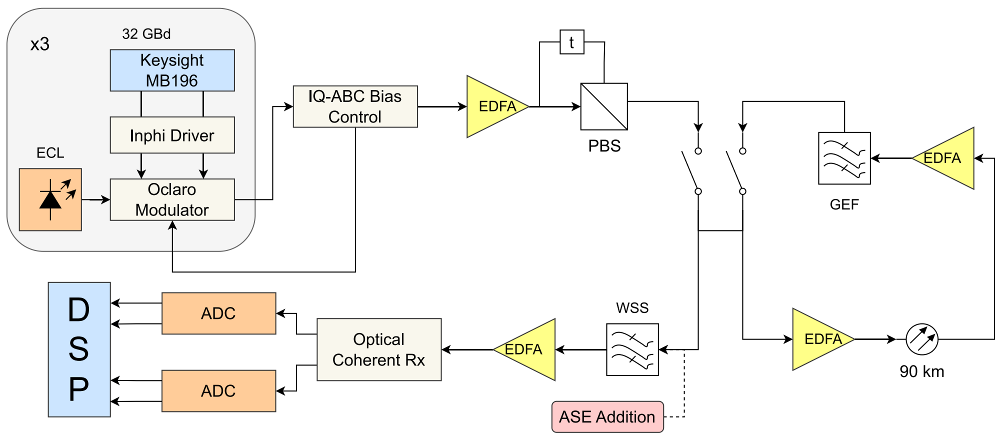

(b) 3 × 32-GBd 16-QAM Nyquist-WDM system (ES2).

Fig. 2. Experimental Nyquist-WDM systems.

and 9 dBm. At the receiving end, a wideband coherent receiver with a 40 GHz bandwidth was employed to detect the optical carriers, with the central carrier serving as the channel of interest.

The received signal is digitized using a real-time oscilloscope with a sampling rate of 80 GSa/s, with each of its four inputs (I and Q in the X and Y polarizations) offering a 35 GHz bandwidth. Offline digital signal processing is performed using MATLAB®, while the proposed DL-based tool is implemented in Python.

The data obtained from both experimental setups, ES1 and ES2, were transformed into IQ diagrams, each representing a time window of approximately 10k symbols. Specifically, ES2 generates three distinct datasets: (i) 0 km with 0 dBm, (ii) 270 km with 0 dB, and (iii) 270 km with 9 dBm. In contrast, ES1 produces only a single dataset from a B2B setup.

3. Development of the monitoring tool for detecting and estimating OSNR and chsp

The proposed monitoring tool operates on grayscale images derived from 2D histograms of the received IQ constellation, where the bin count determines both the spatial resolution of the resulting image and the granularity of the frequency distribution it encodes. In this work, a bin count of 32 is adopted for all reported experiments, as it represents the most favorable trade-off between feature quality and computational efficiency. A brief characterization of performance and training/inference cost across bin configurations of 16, 32, 64, and 128 is presented in Section 4, where this design choice is further substantiated. The proposed feature representation captures the statistical distribution of the received symbols, which may include distortions arising from multiple impairments. These include linear ICI, ASE noise, and potential nonlinear effects. However, the framework does not explicitly isolate these contributions, but rather learns their combined impact on the symbol distribution. Besides, this framework is not restricted to a specific modulation format. Since it relies on statistical and geometric properties of the received constellations, it can be extended to other M-QAM formats, provided that appropriate training data is available. Note that 16-QAM was used in this study as a representative intermediate format sensitive to both noise and interference. On the other hand, the selection of the evaluated architectures is motivated by the need to analyze different levels of feature representation and model complexity within the context of optical performance monitoring. The Fully Connected Neural Network (FCNN) architecture is considered a baseline deep learning model operating on flattened input data, without explicitly exploiting spatial structure. The Autoencoder (AE)-FCNN is included to investigate whether compressed latent representations can improve performance by reducing redundancy in the input features. In contrast, the Convolutional Neural Network (CNN) is specifically designed to capture spatial correlations in the 2D histogram images, making it particularly suitable for this task. In addition to deep learning models, a support vector machine (SVM) is considered as a classical machine learning benchmark, enabling comparison with non-deep-learning approaches commonly used in optical performance monitoring.

3.1. Image preprocessing

Fig. 3 illustrates the step-by-step process of the image preprocessing using a Gaussian filter. After obtaining 2D histograms with 32 values per bin, which clearly reveal the frequency distribution, 32 × 32-pixel

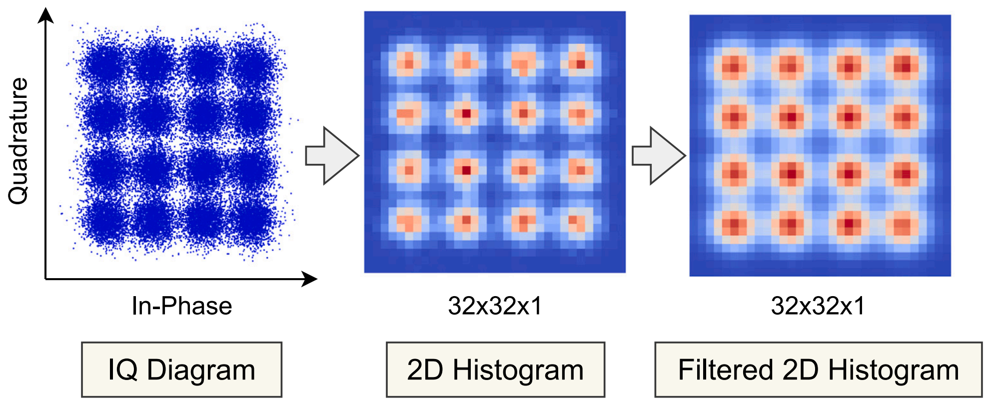

Fig. 3. Step-by-step of obtaining and filtering the 2D histogram from the IQ diagram of the received symbols using a Gaussian filter.

images are generated for each 2D histogram. Then, the noisy image is filtered using a Gaussian filter of size $k \times k$ according to

$$G_i = \alpha \exp\left(\frac{-\left(i - \frac{ksize-1}{2}\right)^2}{2\sigma_X\sigma_Y}\right),$$

where $ksize$ is the kernel size, and $i = 1, \ldots, ksize - 1$. It is necessary to specify the kernel's width and height, which must be positive and odd. Additionally, the standard deviation in the X and Y directions, labeled as σ_X and σ_Y, respectively, must be defined. Commonly, σ_X and σ_Y have the same value. α is a scale factor. If α is non-positive, it is computed from $ksize$ as

$$\alpha = 0.3\,(0.5\,(ksize - 1) - 1) + 0.8.$$

For the Gaussian filter, the sum of all filter coefficients, defined as

$$\sum_{i=1}^{ksize-1} G_i = 1,$$

must be satisfied. The Gaussian blur is highly effective for removing Gaussian noise from an image, making it particularly useful in this approach as it reduces the effects of AWGN on the constellations, thereby emphasizing residual effects that may be characteristic of the ICI remaining in the image, such as the temporal correlations that occur when transmitting overlapping channels. In our previous work [13], we demonstrated that a 5×5-kernel works well for these datasets because a 3×3-kernel does not remove enough Gaussian noise from the image, and kernels larger than 5×5 remove too much valuable information.

After applying the Gaussian filter, the grayscale images, which initially ranged from 0 to 255, now exhibit varying value ranges. To standardize these images and ensure consistency in the dataset, we use a min–max scaler, as

$$X_{scaled} = \frac{X - X_{min}}{X_{max} - X_{min}},$$

where X represents the original pixel intensity, and X_{min} and X_{max} denote the minimum and maximum intensity values observed in the input image, respectively. This scaler rescales pixel values to a uniform range, specifically 0 to 1. The normalization process accelerates network convergence because gradient-based optimization methods, such as stochastic gradient descent or Adam, perform more effectively within a small and controlled range. In this work, we have normalized each image in the training set to avoid leaking information to the test set model. In this sense, the test set is normalized based on the train set's min and max values.

The final dataset used in the proposed models is described as follows: the number of images per scenario varies according to the number of transmitted symbols in each raw dataset: the 16-GBd B2B scenario comprises 729 images, the 32-GBd B2B scenario 2166 images, and the 32-GBd transmission scenarios at 0 dBm and 9 dBm contain 1101 and 1554 images, respectively. Data partitioning follows a stratified 90–10 train–test split per fold, ensuring proportional class representation across folds. For binary classification of ChSp, class 0 corresponds to non-overlapping configurations (ChSp>ChBw) and class 1 to overlapping configurations (ChSp≤ChBw), where ChBw is 17.6 GHz and 35.2 GHz for the 16-GBd and 32-GBd systems, respectively. For binary OSNR classification, class 0 denotes high-OSNR conditions (OSNR> 23 dB) and class 1 denotes low-OSNR conditions (OSNR≤23 dB). This decision threshold is configurable depending on system design criteria; in this work, a value of 23 dB is selected as a representative operating point for evaluation. In multiclass settings, each discrete value of ChSp or OSNR within the ranges described in Section 2 constitutes an independent class.

3.2. Fully connected neural network (FCNN)

An FCNN, also known as a dense neural network, is one of the fundamental architectures used in DL. In an FCNN, every neuron in a given layer is connected to all the neurons in the subsequent layer. This connectivity enables the network to capture complex, non-linear relationships between input and output features. The FCNN consists of an input layer that receives raw data. In this case, 32 × 32 single-channel (grayscale) pixel images are received. These images must be flattened into a vector of 1024 values, where each value corresponds to a pixel from the original image. The network also includes hidden layers composed of multiple layers of neurons. Here, non-linear transformations are applied using activation functions, enabling the network to learn intricate patterns within the data. Finally, the output layer produces the final prediction, which can be used for classification or regression, depending on the nature of the target variable.

We specifically propose an FCNN architecture and a detailed parameterization in Tables 2 and 3, respectively, using an input layer of 1024 neurons, obtained by flattening a 32×32-pixel image. Additionally, we employ two hidden layers with 200 and 100 neurons, utilizing rectified linear unit (ReLU) activation functions. ReLU was chosen because the images are normalized to the range of 0 to 1, with the function defined as x for $x \geq 0$ and 0 for $x < 0$. The output layer of the FCNN varies depending on the specific task, whether it be classification or regression. For classification tasks, the number of neurons in the output layer equals the number of classes that need to be identified, such as different ChSp or OSNR ranges. Cross-entropy is the loss function used for classification, with a sigmoid activation function for binary classification tasks and a SoftMax activation function for multiclass tasks. In contrast, the output layer for regression tasks consists of a single neuron that employs an L1 loss function, using either no activation function or a linear activation function.

Table 2
FCNN architecture and hyperparameters for classification (anomaly detection) and regression (estimation) tasks.

Layer	No. of neurons	Activation function
Input	1024	–
Hidden layer 1	200	ReLU
Hidden layer 2	100	ReLU
Output (Classification)	Number of classes	Sigmoid (binary)/ SoftMax (multiclass)
Output (Regression)	1	Linear

Table 3
Hyperparameter selection.

Hyperparameters	
Batch size	16
Optimizer	Adam
Learning rate	0.001
Patience	30
Loss function	Cross Entropy (classification)/ L1 (regression)
Cross validation	10-Folds
Weight initialization	Random

3.3. Autoencoder for feature extraction and FCNN (AE-FCNN)

An autoencoder (AE) is an unsupervised neural network architecture designed to learn efficient, compressed representations of input data through the process of encoding and decoding. The AE consists of two main components: the encoder and the decoder. Firstly, the encoder transforms high-dimensional input data, such as 32×32-pixel grayscale images, into a lower-dimensional latent space representation, effectively capturing the most relevant features while reducing redundancy. This latent space, often referred to as the bottleneck layer, represents a compressed version of the original data. Finally, the decoder then attempts to reconstruct the original input from this compressed representation, ensuring that the encoded features retain the critical information necessary for accurate reconstruction.

In this work, the autoencoder is trained to minimize reconstruction loss, using MAE as the loss function due to the continuous nature of pixel values. The architecture of AE consists of an encoder compressing the 1024-dimensional input into a latent space ranging from 16 to 128, and the decoder mirrors this structure to reconstruct the input. After training the AE, the encoder component is isolated to serve as a feature extractor. The latent space representations generated by the encoder are then used as inputs to the FCNN described previously. This approach investigates whether these compressed, high-level feature representations contribute to improved performance in estimation and detection tasks compared to raw pixel data. By leveraging the AE's ability to capture essential data features, we hypothesize that the FCNN may achieve enhanced accuracy and robustness, particularly in scenarios with complex patterns or noisy data.

3.4. Convolutional neural network (CNN)

A CNN is a DL architecture designed to process data with a grid-like structure, such as images. Unlike FCNNs, where each neuron is linked to all neurons in the following layer, CNNs utilize local connectivity through convolutional operations. CNNs are inspired by the human visual system, where neurons respond to stimuli within specific receptive fields. This architectural design allows the network to automatically learn hierarchical spatial features by applying learnable filters (kernels) that slide across the input data to identify local patterns. As these local patterns are combined in deeper layers, the network captures more complex and abstract representations, making CNNs particularly adept at tasks where the spatial relationships between pixels are critical.

Table 4
CNN architecture and hyperparameters for classification (anomaly detection) and regression (estimation) tasks.

Layer	No. of filters/neurons	Activation function
Input	32×32-pixel image	–
Conv. 2D layer 1	16	ReLU
MaxPooling layer	–	–
Conv. 2D Layer 2	32	ReLU
Flatten	–	–
Hidden layer 1	200	ReLU
Hidden layer 2	100	ReLU
Output (Classification)	Number of classes	Sigmoid (binary)/ SoftMax (multiclass)
Output (Regression)	1	Linear

In this work, the implemented CNN architecture, detailed in Table 4, comprises two convolutional layers containing 16 and 32 filters, each followed by a ReLU activation function. The convolutional layers apply filters that scan the image, enabling the capture of local features such as data clusters and their scattering due to ICI and ASE noise effects. The weights of the filters are determined by the network's training process and operate hierarchically, meaning that deeper layers tend to capture more abstract features. Meanwhile, the pooling layers discard less relevant features or redundant information. This hierarchical feature extraction process enables the network to effectively represent both local and global patterns in 32×32-pixel grayscale images. Following the convolutional layers, the resulting feature maps are flattened into a one-dimensional vector, which is then input into the fully connected layers. These layers mirror the configuration of a previously described FCNN, consisting of two hidden layers with 200 and 100 neurons, respectively, employing ReLU activation functions. In addition, the hyperparameters of the network are the same as those of FCNN, as shown in Table 3.

This approach enables the network to merge the spatial features extracted by the convolutional layers with the dense representational capabilities of the fully connected layers, thus facilitating the learning of complex and non-linear relationships within the data. The architecture aims to assess whether the convolutional layers' capacity to extract spatial features can enhance performance in estimation and detection tasks compared to models that depend solely on raw pixel data or compressed representations generated by autoencoders.

4. Results and discussion

The proposed monitoring tool for estimating and detecting both OSNR and ChSp is based on three DL architectures: (i) FCNN, (ii) AE-FCNN, and (iii) CNN, as detailed in Sections 3.2, 3.3, and 3.4, respectively. We use accuracy, precision, recall, and F1-score metrics to evaluate performance in classification tasks, such as anomaly detection. For regression tasks involving parameter estimation, such as ChSp and OSNR, we use the mean absolute error (MAE) and the R-squared (R^2) metrics. Confidence intervals or standard deviation across folds were not reported in this work because the metrics used (accuracy and F1-score for classification tasks, and MAE and R^2 for regression tasks) are computed by aggregating predictions over all folds, such that the entire dataset serves as test data exactly once and the decision threshold was fixed to ensure the same class weight, providing an unbiased estimation.

4.1. Anomaly detection results

We evaluated the performance of the proposed DL architectures compared to the classical ML method, SVM, using several metrics, including accuracy, precision, recall, and F1 score for classification tasks. We are mainly interested in analyzing precision and recall metrics because they provide information on the percentage of predictions

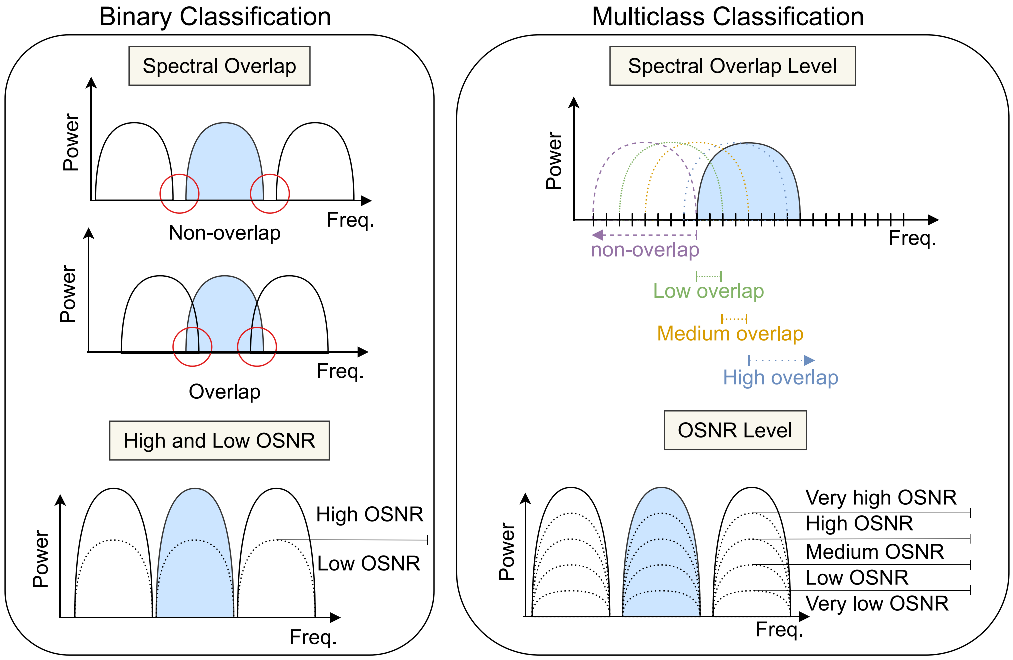

Fig. 4. Representation of the spectral overlap and OSNR classification problem in the spectrum of the three fiber channels, including anomaly detection issues (binary classification) and classification of different levels (multiclass classification).

that the tool classifies as false positives and false negatives. In our problem, false positives refer to scenarios in which the channels are not overlapping (or have a low OSNR) but are erroneously classified as overlapping (or having a high OSNR). On the other hand, false negatives are scenarios in which there is overlap (or a high level of OSNR), but they are incorrectly classified as having no overlap (or a low level of OSNR). This is important because it is necessary to know when the system is genuinely failing in order to make corrections or take preventive measures. Therefore, it is preferable to have more false positives (detecting a failure and taking action) than false negatives (overlooking damage to the system).

These metrics are consistent across the 16- and 32-GBd B2B setups and the 32-GBd setup involving transmission over 270 km at 0 and 9 dBm, providing valuable insights into the models' performance. We classify the dataset based on spectral separation or OSNR to evaluate performance, depending on the problem being addressed. The remaining information is utilized as an additional feature for multimodal models. This means the model has an additional input, in addition to the image (built from constellation diagrams). Fig. 4 illustrates three optical channels for each in the binary and multiclass classification problems for ChSp and OSNR. In the binary problem, the goal is to detect whether there is overlap between channels or whether the OSNR of the channel under evaluation is low. On the other hand, in the multiclass problem, the range of ChSp or OSNR is divided, increasing the granularity of detection, e.g., high, low, and no spectral overlap.

Fig. 5 presents the anomaly detection results for the 1-6GBd B2B, 32-GBd B2B, and 32-GBd with 270 km of fiber at 0 and 9 dBm setups proposed in this study. In the 16-GBd setup, all DL-based detection models achieved approximately 98% accuracy and F1-score in spectral overlap detection (binary classification), while SVM achieved around 96% [see Figs. 5(a) and 5(b)]. These results indicate a balanced identification of true positives and a high proportion of correct positive predictions, even without including OSNR information as an additional feature. For detecting different spectral overlap ranges in the same 16-GBd setup (multiclass classification), accuracy metrics ranged between 61% and 75%, while F1-score exceeded 80%, demonstrating the robustness of the models in handling class imbalance [see Figs. 5(c) and 5(d)]. On the other hand, in the 32-GBd setups, which include B2B and 270 km of fiber at 0 and 9 dBm, performance significantly decreased compared to the 16-GBd setup. In binary classification, the F1-score was slightly higher than the accuracy, both with and without OSNR information, reaching values between 88% and 91%. However, this does not fully reflect the actual behavior of the phenomenon, as precision and recall metrics show considerable differences when the OSNR value is not included: recall exceeds 90%, while precision falls below 75%. This result indicates that DL architectures effectively identify most actual spectral overlap cases (true positives) but also misclassify many cases as overlapping when they are not (false positives). Moreover, for multiclass classification, performance reached approximately 50% when OSNR information was included, but dropped to 30% for DL architectures and 16% for SVM when this additional feature was not included. The results obtained for the 270 km transmission at 9 dBm show a slight improvement in performance compared to the 0 dBm case, even under equivalent OSNR conditions. This indicates that the gain is not driven by noise levels, but by the nature of the distortions at higher launch powers. In particular, nonlinear effects introduce more structured patterns in the received constellations, which are effectively captured by the 2D histogram representations, enhancing feature separability. As a result, the proposed model achieves improved discrimination between channel conditions, despite the coexistence of nonlinear impairments with linear ICI and ASE noise. It should be noted that the framework captures the combined impact of these impairments at the feature level, rather than explicitly separating their individual contributions

Similarly, in Fig. 6, when identifying whether the OSNR level in the system is high or low, all models achieved over 96% accuracy and F1-score values between 93% and 99% in all configurations when ChSp information was included. However, when the additional feature was not provided, performance in the 16-GBd setup remained largely unaffected, whereas in the 32-GBd setups, accuracy dropped to 81%–85%, with similar F1-score values. Notably, the 32-GBd setup at 0 dBm was more affected, likely because this configuration is more sensitive to noise, making it difficult to distinguish between signal and noise levels effectively [see Fig. 6(a) and 6(b)]. Regarding detecting different noise levels, when ChSp information was included, results ranged from 80% to 90% in accuracy and F1-score. However, if the tool did not include

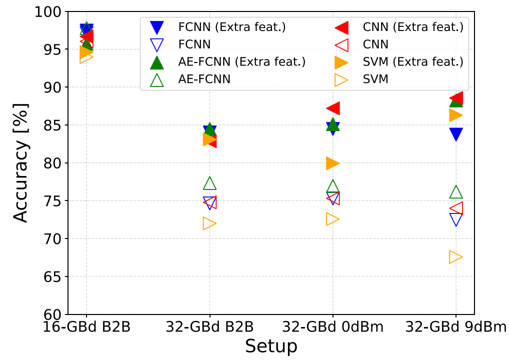

(a) Accuracy in spectral overlap detection.

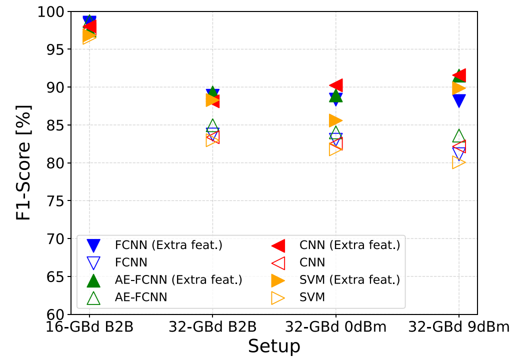

(b) F1-Score in spectral overlap detection.

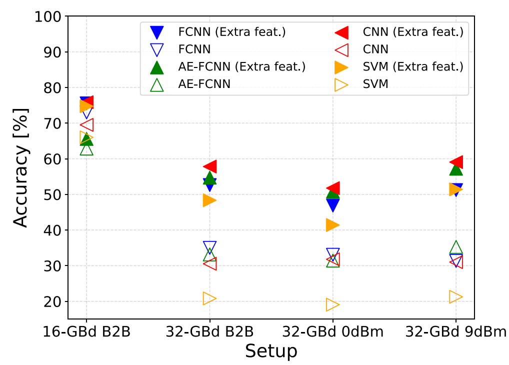

(c) Accuracy in the detection of overlapping spectral ranges.

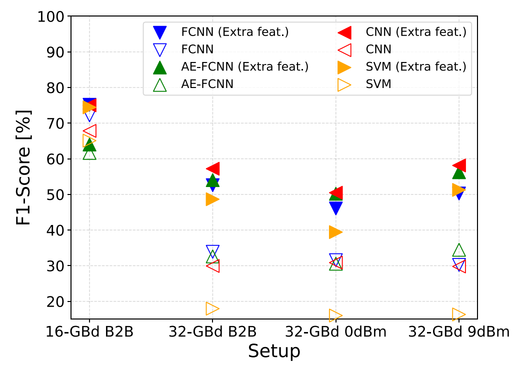

(d) F1-Score in the detection of overlapping spectral ranges.

Fig. 5. Spectral overlap detection results for the four setups with four different models, including and not including OSNR information to the classifier.

this information, performance dropped below 60% in accuracy and F1-score, falling below 50% in the 32-GBd setups when using SVM. This result indicates the model struggles to make accurate predictions when it lacks sufficient system information [see Fig. 6(c) and 6(d)]. Unlike the spectral overlap detection problem, the results of the proposed DL architectures maintain a balance between false positives and false negatives (equal recall and precision), indicating that they are not biased toward one type of error, but rather make mistakes in equal proportions. It is important to note that this multiclass scenario represents a fine-grained classification problem under extreme spectral overlap conditions, where class boundaries are inherently ambiguous due to the similarity between impairment patterns. As such, the observed performance reflects the intrinsic difficulty of the task rather than a limitation of the proposed approach. From an operational perspective, anomaly detection is typically addressed through coarser and more actionable indicators, where higher performance is achieved. Complementarily, the estimation of continuous parameters such as spectral separation and OSNR is addressed in the following subsection, where the proposed method demonstrates accurate predictive capability.

To further complement the evaluation in binary classification scenarios, ROC–AUC curves are presented for the FCNN model, as shown in 7(a) and 7(b). These correspond to low-OSNR detection and spectral overlap detection, respectively, and provide a threshold-independent assessment of classifier performance in terms of the trade-off between True Positive Rate (TPR) and False Positive Rate (FPR). The observed behavior is consistent with the trends indicated by the weighted F1-score, confirming that the proposed approach achieves reliable discrimination capability under both impairment conditions.

A comparative analysis of the ROC curves across operating conditions reveals that the highest physical signal integrity (e.g., B2B operation) does not necessarily correspond to the best classification performance. In particular, increasing the symbol rate from 16- to 32-GBd leads to a degradation in AUC, which can be associated with increased inter-symbol interference and spectral overlap effects.

Interestingly, within the 32-GBd regime, the scenario with higher launch power (9 dBm) exhibits improved classification performance compared to the linear regime (0 dBm). This behavior suggests that nonlinear distortions may introduce additional structure in the received signal that can be exploited by the learning model. While a detailed physical interpretation is beyond the scope of this work, these results indicate that the proposed approach is capable of capturing discriminative features even under nonlinear propagation conditions.

4.2. Parameter estimation results

Parameter estimation aims to measure the spectral separation between the channel under evaluation and its adjacent channels, or to separate the contribution of ASE noise from the signal power, excluding the contribution of adjacent channels. For these estimation tasks, we evaluated the performance of the proposed DL architectures using MAE and R^2 metrics. Similar to anomaly detection tasks (classification), the metrics remain consistent across all setups. Additionally, we assessed the performance of models that do not have access to OSNR or ChSp information (depending on the specific problem). We compared them to models that include this information in the regressor. Fig. 8 illustrates the objective of estimating ChSp and OSNR in the spectrum of optical fiber channels. Specifically, MAE measures the estimation error in GHz or dB, depending on the context of the study.

Fig. 9 presents the spectral spacing estimation results for the three proposed DL architectures across the 16-GBd B2B, 32-GBd B2B, and 32-GBd with 270 km of fiber at 0 and 9 dBm setups. Generally, an MAE below 2 GHz was achieved, with the 16-GBd setup reaching

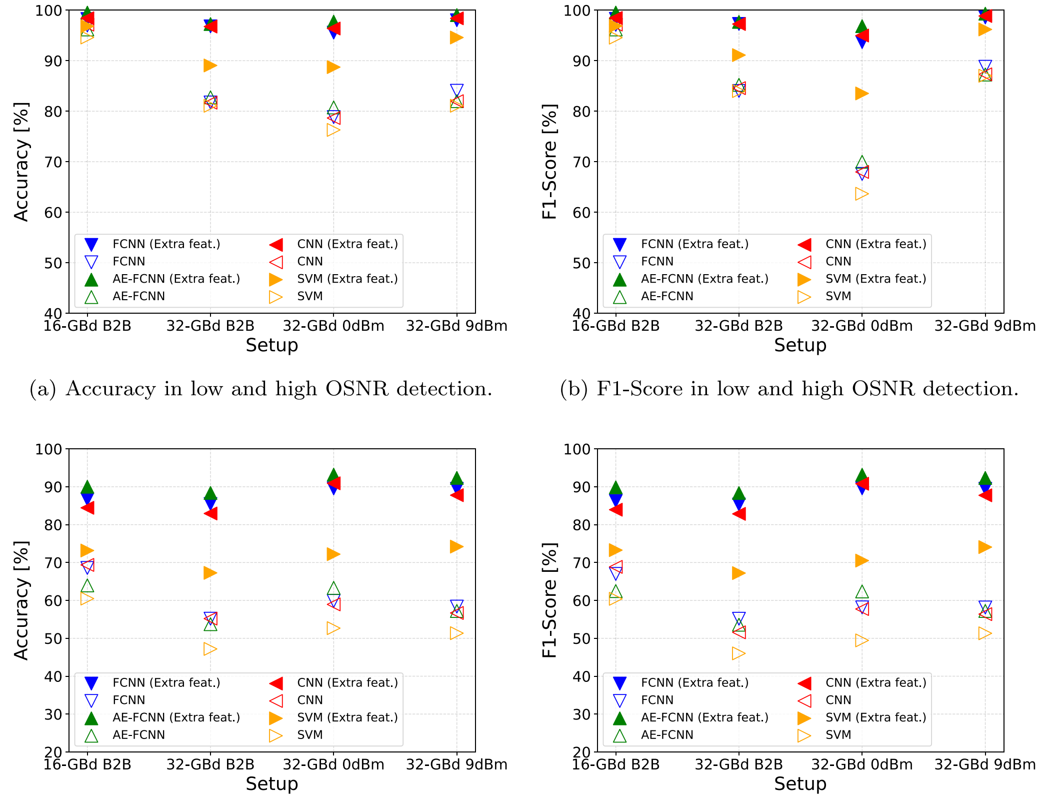

(a) Accuracy in low and high OSNR detection.
(b) F1-Score in low and high OSNR detection.
(c) Accuracy in the detection of OSNR ranges.
(d) F1-Score in the detection of OSNR ranges.

Fig. 6. Low and high OSNR level-detection results for the four setups with four different models, including and not including spectral spacing information to the classifier.

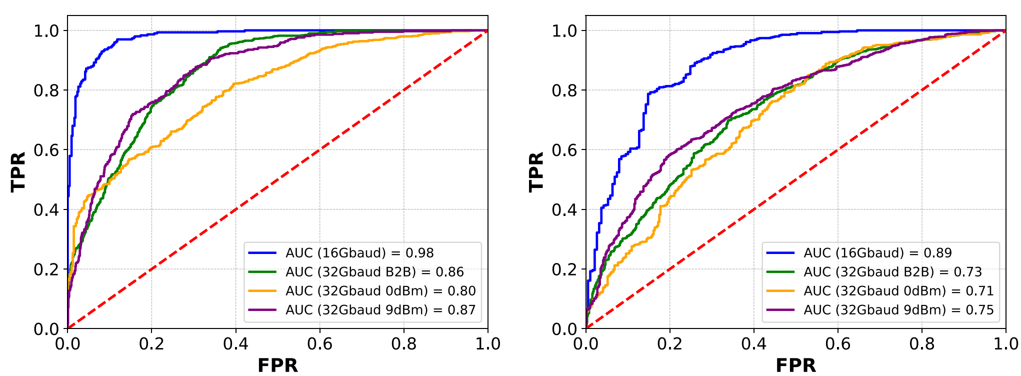

(a) Low-OSNR level-detection curves for the FCNN architecture.
(b) Spectral overlap detection curves for the FCNN architecture.

Fig. 7. ROC–AUC curves for low-OSNR and spectral overlap detection results for the four setups with the FCNN architecture.

an MAE below 0.5 GHz, demonstrating that it is the most straightforward scenario to model. Furthermore, including OSNR information significantly improved model performance, reducing the MAE to below 1 GHz in the 32-GBd B2B setup, compared to approximately 2 GHz when OSNR information was not included [see Fig. 9(a)]. In terms of R^2, only the CNN achieved values above 0.6 in the 32-GBd setups, whereas in the 16-GBd setup, all models reached values above 0.8. This suggests that the 32-GBd setups are more challenging to model [see Fig. 9(b)]. Additionally, the CNN outperformed the other models in spectral spacing estimation, achieving lower MAE values and higher R^2 scores than the rest. At the same time, FCNN obtained the worst results, even when the additional feature was included. On the other hand, SVM, as a classical ML algorithm, outperformed both FCNN and AE-FCNN in most scenarios. The reasons underlying these performance differences across all evaluated models are discussed at the end of this subsection.

Regarding OSNR estimation, Fig. 10 shows that CNN once again delivers the best performance in terms of MAE, achieving errors below 0.5 dB for the 16-GBd scenario and a range of 0.8–1.6 dB for all the 32-GBd scenarios when ChSp information was included. SVM follows, demonstrating better performance in the 32-GBd setups than the 16-GBd setup. The maximum error obtained by the proposed DL architectures, including SVM, was approximately 2.5 dB when ChSp information was not included. CNN remains the best model for OSNR estimation, achieving an R^2 value of approximately 0.9. Again, SVM outperforms FCNN and AE-FCNN, demonstrating better adaptability

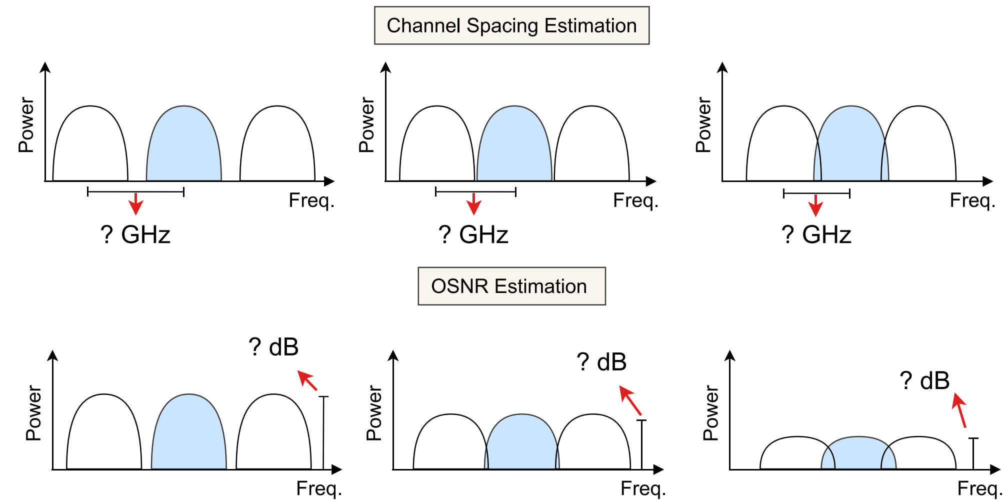

Fig. 8. Sketch of three optical channels representing the problem of estimating the OSNR level and ChSp.

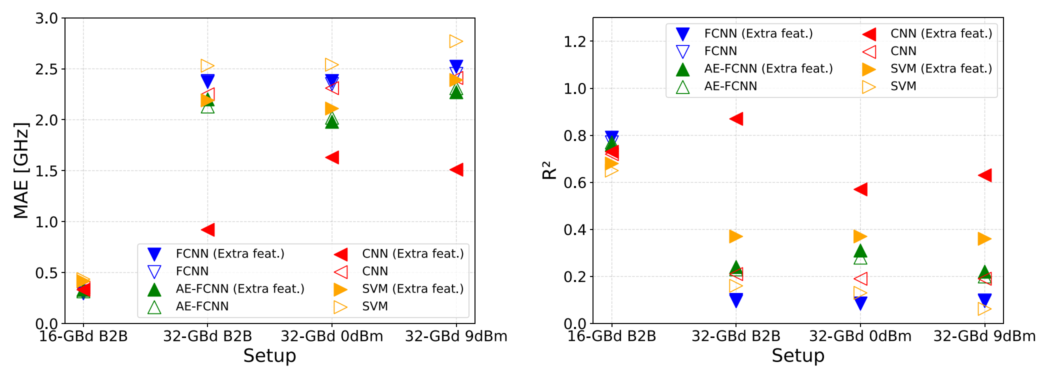

(a) MAE in spectral overlap estimation.

(b) R^2 in spectral overlap estimation.

Fig. 9. Spectral overlap estimation results for the four setups with four different models, including and not including OSNR information to the regressor.

to the dataset. FCNN and AE-FCNN capture only 0.5–0.6 of the data variability. However, this is significantly higher than in spectral spacing estimation, where these models failed to capture even 0.4 of the total variability. This demonstrates that CNN outperforms the other models in parameter estimation when the additional feature is available. However, its performance is similar to that of other models when this information is excluded, highlighting the importance of this feature for reliably estimating system parameters. Across both estimation and anomaly detection tasks, a consistent performance ranking was observed among the evaluated models. The superior performance of the CNN is attributed to its ability to exploit spatial correlations in the 2D histogram representation through convolutional filters, which are not captured by FCNN-based models operating on flattened inputs. The AE-FCNN, despite its ability to reduce input redundancy through unsupervised compression, introduces a bottleneck stage that may discard fine-grained discriminative information, particularly in scenarios where impairment patterns are subtle or closely spaced. Regarding SVM, its competitive performance relative to FCNN and AE-FCNN is consistent with the relatively low-dimensional and structured nature of the input features: the ϵ-tube tolerance margin acts as an implicit regularization mechanism, providing robustness against overfitting that FCNN-based models partially lose due to premature early stopping. Nevertheless, CNN-based models provide a more flexible and extensible framework for capturing spatial patterns and are better suited for generalization to more complex monitoring scenarios.

4.3. Deployment feasibility

The practical deployment of the proposed monitoring framework in elastic optical networks is facilitated by its operation at the symbol level. The method relies on IQ constellation data already available after the DSP stage in coherent receivers, thereby avoiding the need for modifications to the receiver architecture. The feature extraction process, based on 2D histograms, can be implemented as a lightweight preprocessing step, while the trained models can be executed as an auxiliary monitoring module. This module can be integrated either locally within the receiver hardware or at a higher network management layer.

The computational feasibility of this deployment is supported by the training and inference time measurements reported in Tables 5 and 6, which characterize the FCNN and CNN architectures across different bin sizes (16, 32, 64, and 128), evaluated on the 32-GBd scenarios at 0 and 9 dBm, conducted on an NVIDIA GeForce RTX 4070 Ti GPU. Note that while a 32-bin configuration was selected for all experiments reported in this work, the timing results across different bin sizes are provided to illustrate the computational trade-off associated with increasing input resolution. For the CNN, training time increases substantially with bin count, rising from approximately 0.8 s at 16 bins to over 11 s at 128 bins per 100 epochs. For the FCNN, the increase is more moderate, ranging from approximately 0.6 s to 2.4 s under the same conditions. Inference times remain below 0.013 s in all cases, confirming compatibility with near-real-time operation on standard processing units.

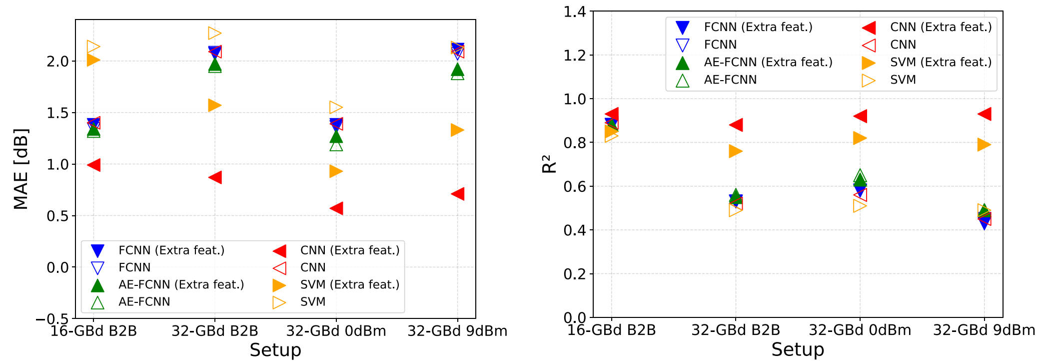

(a) MAE in OSNR estimation. (b) R^2 in OSNR estimation.

Fig. 10. OSNR estimation results for the four setups with four different models, including and not including spectral spacing information to the regressor.

Table 5
Characterization of the FCNN through the number of bins for training and inference.

Bins	32-GBd 0 dBm		32-GBd 9 dBm	
	Training [s]/ 100 epochs	Inference [s]	Training [s]/ 100 epochs	Inference [s]
16	0,5567	0,0003	0,8004	0,0003
32	0,5953	0,0003	0,8760	0,0004
64	0,8813	0,0008	1,2854	0,0011
128	2,4445	0,0027	3,3894	0,0040

Table 6
Characterization of the CNN through the number of bins for training and inference.

Bins	32-GBd 0 dBm		32-GBd 9 dBm	
	Training [s]/ 100 epochs	Inference [s]	Training [s]/ 100 epochs	Inference [s]
16	0,8463	0,0003	1,1002	0,0013
32	0,8236	0,0004	1,1546	0,0016
64	1,9397	0,0027	2,7224	0,0034
128	11,3134	0,0089	15,9215	0,0125

Although the experiments are conducted using a three-channel Nyquist-WDM configuration, the proposed methodology is not limited to this setup. The monitoring framework operates on a per-channel basis, where feature extraction and inference depend on the local interference conditions rather than the total number of channels in the system. In dense WDM scenarios, the dominant impairments affecting the channel of interest arise primarily from adjacent channels, suggesting that the proposed approach can generalize to higher channel counts without structural modifications. Regarding modulation formats, the methodology relies on statistical and geometric properties of the received constellations and can be extended to other M-QAM formats through retraining with appropriate datasets. Experimental validation under fully loaded systems and additional modulation formats constitutes a relevant direction for future work.

5. Conclusion

We proposed and experimentally validated a DL-based OPM method for elastic WDM networks, addressing the tasks of anomaly detection and parameter estimation (OSNR and ChSp) through a task-specific configuration of three distinct neural network architectures, employing 2D histogram analysis. The introduced architectures (FCNN, AE-FCNN, and CNN) were configured individually for classification or regression, depending on the target monitoring function. When operated in classification mode, the proposed models effectively identified anomalies associated with linear ICI and OSNR degradation due to ASE noise, enabling detection of conditions such as spectral overlap or low OSNR. Conversely, when configured for regression tasks, the same architectural frameworks facilitated accurate estimation of ChSp between the evaluated channel and its neighboring channels, as well as the corresponding signal power levels. The results demonstrate that the proposed approach is robust against linear ICI effects and facilitate the differentiation between ICI and ASE noise effects while operating in a blind manner, without requiring prior knowledge of system parameters or auxiliary signaling. Although blind operation inherently incurs a bounded performance penalty, estimation errors remained below 2 dB for OSNR and 2.5 GHz for ChSp, while errors below 1 dB and 1.5 GHz were obtained during prior system characterization. In addition, the method accurately identified anomalies such as spectral overlap and low-OSNR events, achieving detection accuracies above 95% for the 16-GBd scenario in both detection tasks, while for 32-GBd transmission, anomaly-detection accuracies up to 85% and 99% were obtained for spectral overlap and low-OSNR detection tasks, respectively. This reduction is mainly attributed to the increased spectral occupancy at higher symbol rates, which intensifies channel overlap and reduces the separability between linear ICI and ASE noise in the symbol domain. On the other hand, the presence of nonlinear effects in transmission scenarios is implicitly captured within the proposed feature representation, although their individual contributions are not explicitly isolated. Among the evaluated models, the CNN architecture exhibited greater reliability than FCNN-based solutions, owing to its enhanced ability to extract spatial features from 2D histogram images. The practicality of the proposed method is further supported by its blind detection and estimation, which are based exclusively on symbol-level information, thereby simplifying deployment in real-world elastic optical networks. The per-channel operation and format-agnostic feature representation of the proposed framework support its scalability to denser and more heterogeneous network scenarios beyond the three-channel 16-QAM configuration used for validation. The proposed framework can be deployed as a low-complexity monitoring module integrated into DSP-assisted receivers or network management systems, facilitating its adoption in ROADM-based elastic optical networks. The relatively low computational complexity of the proposed models, evidenced by inference times below 0.013 s, supports their potential implementation in real-time monitoring applications, although hardware-specific benchmarking on embedded platforms is left for future work. Overall, these findings position the proposed monitoring framework as a flexible, adaptable alternative to conventional optical performance monitoring tools, thereby enabling enhanced automation and robustness in future elastic WDM systems.

CRediT authorship contribution statement

Kevin David Martinez-Zapata: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Jhon James Granada-Torres:** Writing – review &

editing, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the following project from the University of Antioquia: ''Prototipo de red óptica de acceso usando fuentes ópticas programables de grafeno *y* Machine Learning [2024-76210]''. Special thanks are extended to Prof. Stephen E. Ralph, Director of the Georgia Electronic Design Center, for his support and institutional encouragement.

Data availability

Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

References

[1] W.S. Saif, A.M. Ragheb, B. Nebendahl, T. Alshawi, M. Marey, S.A. Alshebeili, Machine learning-based optical performance monitoring for super-channel optical networks, in: Photonics, vol. 9, MDPI, 2022, p. 299.

[2] Z. Wan, Z. Yu, L. Shu, Y. Zhao, H. Zhang, K. Xu, Intelligent optical performance monitor using multi-task learning based artificial neural network, Opt. Express 27 (8) (2019) 11281–11291.

[3] A. Hraghi, L. Minelli, A. Nespola, S. Piciaccia, G. Bosco, Experimental demonstration of linear inter-channel interference estimation based on neural networks, IEEE Photonics J. 15 (2) (2023) 1–6.

[4] J. Thrane, J. Wass, M. Piels, J.C. Diniz, R. Jones, D. Zibar, Machine learning techniques for optical performance monitoring from directly detected PDM-QAM signals, J. Lightwave Technol. 35 (4) (2016) 868–875.

[5] K. Kaeval, F. Slyne, S. Troia, E. Kenny, K. Große, H. Griesser, D.C. Kilper, M. Ruffini, J.-J. Pedreno-Manresa, S.K. Patri, G. Jervan, Employing channel probing to derive end-of-life service margins for optical spectrum services. To appear in OPTICA journal of optical communications and networking, 2023, arXiv:2302.04623.

[6] D. Zhang, H. Zhang, Y. Cao, Y.L. Xue, Combined modulation format identification and optical signal-to-noise ratio monitoring with high accuracy and generalizability based on a proposed fused module few-shot learning algorithm in dynamic coherent optical transmissions, Opt. Express 32 (8) (2024) 14719–14734, http://dx.doi.org/10.1364/OE.511019, URL https://opg.optica.org/oe/abstract.cfm?URI=oe-32-8-14719.

[7] J.J. Granada Torres, S. Varughese, V.A. Thomas, A. Chiuchiarelli, S.E. Ralph, A.M. Cárdenas Soto, N. Guerrero González, Mitigation of time-varying distortions in nyquist-WDM systems using machine learning, Opt. Fiber Technol. 38 (2017) 130–135, http://dx.doi.org/10.1016/j.yofte.2017.09.008, URL https://www.sciencedirect.com/science/article/pii/S106852001730113X.

[8] Y. Zhou, H. Yin, J. Xiong, S. Song, J. Zhu, J. Du, H. Chen, Y. Tang, Overview and performance analysis of various waveforms in high mobility scenarios, in: 2024 7th International Conference on Communication Engineering and Technology, ICCET, 2024, pp. 35–40, http://dx.doi.org/10.1109/ICCET62255.2024.00013.

[9] E. London, E. Virgillito, A. D'Amico, V. Curri, Accumulation of cross-channel non-linear interference in dispersion-managed and disaggregated optical network segment, 2024, http://dx.doi.org/10.48550/arXiv.2402.09415, arXiv:2402.09415.

[10] C.A.A. Anacona, H.I.C. Martínez, C.A.M. Ocampo, J.J.G. Torres, OSNR estimation using IQ histograms and machine learning in gridless WDM systems, in: Frontiers in Optics + Laser Science 2025 (FiO, LS), Optica Publishing Group, 2025, p. JW5A.3, URL https://opg.optica.org/abstract.cfm?URI=LS-2025-JW5A.3.

[11] F. Locatelli, K. Christodoulopoulos, M.S. Moreolo, J.M. Fàbrega, L. Nadal, S. Spadaro, Spectral processing techniques for efficient monitoring in optical networks, J. Opt. Commun. Netw. 13 (7) (2021) 158–168, http://dx.doi.org/10.1364/JOCN.418800, URL https://opg.optica.org/jocn/abstract.cfm?URI=jocn-13-7-158.

[12] K.D.M. Zapata, S.E. Ralph, J.J.G. Torres, IQ constellation images and deep learning-based method for estimating OSNR in gridless WDM systems, in: Integrated Photonics Research, Silicon and Nanophotonics, Optica Publishing Group, 2025, pp. JD1–8.

[13] K.D.M. Zapata, J.J.G. Torres, Estimation of spectral spacing and linear ICI effects using iq constellation images and CNN in gridless WDM systems, Appl. Opt. 64 (3) (2025) 634–640, http://dx.doi.org/10.1364/AO.539743, URL https://opg.optica.org/ao/abstract.cfm?URI=ao-64-3-634.