

Explainability-Driven Comparison of Topology-Aware and Feature-Based QoT Estimation in Multiband Optical Networks

Original

Explainability-Driven Comparison of Topology-Aware and Feature-Based QoT Estimation in Multiband Optical Networks / Eid, J., Galante, F., Rottondi, C., Bianco, A.. - ELETTRONICO. - (2026), pp. 1-6. (2026 IFIP Networking Conference, IFIP Networking 2026 Lugano (CH) 24-27 May 2026) [10.23919/ifipnetworking70592.2026.11579065].

Availability:

This version is available at: 11583/3014554 since: 2026-08-17T08:45:51Z

Publisher:

Institute of Electrical and Electronics Engineers

Published

DOI:10.23919/ifipnetworking70592.2026.11579065

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Explainability-Driven Comparison of Topology-Aware and Feature-Based QoT Estimation in Multiband Optical Networks

Jana Eid, Franco Galante, Cristina Rottondi, Andrea Bianco
Department of Electronics and Telecommunications, Politecnico di Torino, Italy

Abstract—Accurate estimation of the Signal-to-Noise Ratio (SNR) of prospective lightpaths is a fundamental requirement for reliable operation of dynamic multiband optical networks. In recent years, a variety of machine learning approaches have been proposed for SNR prediction, ranging from topology-aware Graph Neural Networks (GNNs) to feature-based ensemble models such as eXtreme Gradient Boosting (XGB). While their predictive performance has been widely evaluated in terms of accuracy-related metrics, a systematic comparison of their decisional logic and results interpretability remain largely unexplored.

This paper focuses on model explainability and offers a post-hoc interpretability study of two predictors, based on GNN and XGB, using Shapley additive explanations to examine feature importance patterns in multiband scenarios. Our analysis shows that, although XGB typically achieves lower aggregate prediction errors, it predominantly relies on few path-level descriptors. In contrast, the GNN distributes predictive influence across topology-aware embeddings and full-spectrum occupancy representations, reflecting a deeper exploitation of structural network information.

I. INTRODUCTION

Quality of Transmission (QoT) estimation is essential for lightpaths (LP) provisioning in optical networks. In dynamic C+L multiband systems, QoT estimation becomes increasingly difficult due to heterogeneous amplification and cross-band interactions [1]. Machine learning (ML) models for QoT estimation have emerged as powerful alternatives to analytical formulations [2], [3], learning intricate relationships between physical-layer parameters, network topology, and spectral state directly from operational data. However, different model families encode network information and LP characterization in fundamentally different ways: some rely on compact, pre-defined path-level features, while others incorporate global topology and structural dependencies.

In [4], we compared a topology-aware Graph Neural Network (GNN) with a feature-based eXtreme Gradient Boosting (XGB) regressor for Signal-to-Noise Ratio (SNR) prediction in a C+L optical network with incremental loading scenario. Results show that the GNN model outperforms XGB in early and sparse spectrum occupation phases, whereas XGB achieves lower aggregate error in heavily-loaded network conditions but exhibits systematic SNR overestimation for long-haul LPs. This highlights a critical limitation of aggregate error metrics: they fail to capture how prediction errors are distributed and

whether they induce operationally unsafe behaviors. Indeed, from an operator's perspective, overestimation of SNR leads to risky provisioning decisions, whereas underestimation results in overly conservative resource usage. As a consequence, evaluating models only through average error metrics does not uncover the underlying decisional mechanisms that ultimately determine their practical usefulness.

This limitation motivates the need for a feature-attribution-level analysis capable of explaining why models behave differently in various operational conditions. In particular, understanding whether a model relies on a narrow subset of features, how it encodes topology and spectral information, and how these dependencies evolve with band occupation and network load cannot be inferred from performance metrics alone. To address this gap, in this paper, we conduct a systematic explainability study based on Shapley additive explanations (SHAP values), decomposing predictions into feature-level contributions and inspecting their variations across time-evolving multiband scenarios. This permits not only to quantify which features drive predictions, but also to uncover fundamental differences in how topology-aware and feature-based models encode physical-layer impairments and adapt to changing spectral loads. In particular, the XGB model consistently focuses on a small subset of path-level features, most notably path length, resulting in a largely invariant decisional logic across traffic loads and spectral bands. In contrast, the GNN exhibits a more adaptive behavior, leveraging topology-aware embeddings and full-spectrum occupancy information to capture load-dependent and band-specific effects.

The remainder of the paper is organized as follows. Sec. II reviews related work. Sec. III describes the model architectures and SHAP evaluation methodology. Sec. IV presents the explainability analysis, and Sec. V concludes the paper.

II. RELATED WORK

ML is a well-established paradigm for QoT estimation, offering data-driven alternatives to analytical models [2], [3], [5]. XGB and gradient boosting regressors have emerged as particularly competitive, offering strong generalization with low inference latency [6].

Although XGB models have proven reliable in accurate QoT estimation, they remain blind to topology and long-range spectral context. GNNs have attracted attention as a topology-aware alternative, capturing multi-hop dependencies inacces-




Fig. 1. Overview of the proposed framework: input representations, GNN and XGBoost model architectures, and post-hoc SHAP explainability pipeline.

sible to flat feature-based models [7]. [8] first demonstrated Graph Convolutional Networks (GCN)-based QoT estimation for unseen network states. [9] showed graph-based learning outperforms per-LP deep neural networks with over 92% accuracy. [10] applied GNNs to joint optical SNR and power estimation. Our previous work [4] first compared GNN and XGB under incremental C+L multiband loading, revealing systematic SNR overestimation in XGB for long-haul LPs, with no explainability analysis.

Regarding the application of eXplainable Artificial Intelligence (XAI) frameworks to optical networking scenarios, [6], [11] applied SHAP to an XGB-based QoT classifier, while [12] demonstrated SHAP-guided feature selection reduces telemetry overhead. [13] applied SHAP to tree-based fault diagnosis. However, a SHAP-based cross-architecture comparison between topology-aware and feature-based QoT estimators has not been conducted, motivating the present analysis of the models proposed in [4].

III. METHODOLOGY

This section describes the two predictive models under comparison and the post-hoc explainability framework used to analyze their decisional logic. We first present the GNN-based predictor (Sec. III-A), which ingests the full network topology and per-link spectral state, then the XGB-based model (Sec. III-B), which operates on a compact handcrafted feature vector. Finally, Sec. III-C introduces the SHAP-based evaluation methodology, including the metrics used to quantify attribution concentration and cross-scenario stability. An overview of the complete pipeline is shown in Fig. 1.

A. GNN-Based Model

In this topology-aware approach, the optical network is modeled as a graph $G = (V, E)$, where vertices represent network nodes and edges represent directed fiber links. To enable link-centric message passing through the GNN, G is first transformed into its corresponding line graph, in which each node represents a physical link of the original topology. Each node in the line graph is then associated with a feature

vector encoding the spectral state of the corresponding network link. In particular, per-link features capture spectral occupation across all frequency slots, including attributes of occupying LPs when present.

The GNN architecture (Fig.1) consists of three main stages:

- **GCN Layers:** Four GCN layers propagate and aggregate information across neighboring links, enabling the model to capture multi-hop topological dependencies and cross-link spectral interactions.
- **Graph pooling:** Mean aggregation is applied over the subset of links traversed by the candidate LP, yielding a path-level embedding that summarizes its structural and spectral context.
- **Readout stage:** A Multi-Layer Perceptron network processes the aggregated embedding, concatenated with descriptors of the candidate LP (i.e., total path length, number of traversed links, length of the longest link, central frequency (in THz), traffic volume, and modulation format), to produce the final SNR prediction.

B. XGB-Based Model

The feature-based approach relies on XGB, trained on handcrafted descriptors summarizing the characteristics of the prospective LP and its spectral environment. The model employs six features: total path length, number of traversed links, length of the longest link, central frequency, traffic volume, and modulation format. Furthermore, a set of descriptors of the closest left and right spectral neighbors is added to incorporate local spectral context. However, the model remains restricted to predefined local features and does not exploit global topology-aware representations.

C. SHAP Post-Hoc Evaluation

To analyze and compare the decisional logic of the two predictors, we compute SHAP values for both models in a post-hoc fashion. For the XGB regressors, we use TreeSHAP, which provides exact Shapley value computation for tree-based models [14]. For the GNN model, where no closed-form SHAP solution is available, we employ a KernelSHAP

TABLE I
DATASET FEATURE RANGES ($N = 16,678$ LIGHTPATHS)

Feature	Min	Max	Unit
Num links	1	7	–
Path length	40	1120	km
Length of longest link	40	320	km
Frequency (C-band)	191.7	196.5	THz
Frequency (L-band)	186.4	191.5	THz
Traffic	100	300	Gbps
SNR	14.65	34.01	dB

approximation to estimate feature attributions. SHAP provides an additive decomposition of a model’s prediction into feature-level contributions [14]: $f(x) = \phi_0 + \sum_{i=1}^M \phi_i(x)$, where ϕ_0 is the baseline prediction and $\phi_i(x)$ is the marginal contribution of feature i . We aggregate these into the mean absolute value $|\phi_i|$ for global ranking and compute the mean signed value $\bar{\phi}_i$ to capture directional tendencies across samples.

We define the *concentration ratio* to quantify attribution concentration

$$C = \frac{\max_j |\phi_j|}{\sum_j |\phi_j|}, \quad (1)$$

which normalises the dominance of the most influential feature by the total attribution mass. This metric provides a compact summary of how evenly the model distributes importance across features: $C \rightarrow 1$ indicates one feature monopolises the explanation, whereas $C \rightarrow 0$ reflects a more distributed attribution.

To further assess structural consistency across model variants, we compute rank-based correlation metrics between their importance vectors $\{|\phi_i|\}_i$ using Spearman rank correlation ρ [15], which measures the monotonic agreement between feature rankings and is particularly robust to nonlinear rescaling of importance values [16].

IV. NUMERICAL RESULTS AND DISCUSSIONS

A. Simulation Scenarios

We leverage the same synthetic dataset adopted in [4], using the simulation framework of [17], which consists of an incremental traffic-loading scenario, where lightpaths are routed on the shortest path via a first-fit spectrum allocation policy, in a WDM multiband (C+L) optical network based on the 14-node topology [17]. Dataset features are summarized in Tab. I. The dataset is partitioned into four epochs: **Epoch A** (C-band, sparse load), **Epoch B** (C-band, denser occupancy), **Epoch C** (C+L transition, with first L-band LPs being allocated), and **Epoch D** (C+L bands, high spectral load). Four train→test scenarios are evaluated: i) A→B, ii) AB→C (C-band only), iii) ABC→D multiband, and iv) two band-specific ABC→D C-band and C→D L-band variants, tested via an Ensemble model.

B. Accuracy Summary

Tab. II reports Mean Absolute Error (MAE) and Mean Square Error (MSE) for all scenarios. The GNN outperforms

TABLE II
PERFORMANCE SUMMARY (MAE AND MSE)

Scenario	MAE [dB]		MSE [dB]	
	GNN	XGB	GNN	XGB
Epoch B (C-band)	0.358	0.453	0.264	0.331
Epoch C (C-band)	0.214	0.202	0.108	0.096
Epoch D (C-band only)	0.228	0.211	0.121	0.103
Epoch D (L-band only)	0.276	0.249	0.158	0.139
Epoch D (C+L multiband)	0.241	0.209	0.133	0.101

XGB in Epoch B, where sparse occupancy and limited training data favor topology-aware generalization. From Epoch C onward, XGB achieves lower aggregate error, with the largest gap in the full C+L multiband scenario. The subsequent SHAP analysis reveals that these aggregate numbers mask structurally meaningful differences in model behavior.

C. SHAP Analysis

For our interpretability analysis we leverage beeswarm representations for each scenario (see Fig. 2), where each row represents a feature ranked by mean absolute SHAP value. Furthermore, we quantify $|\phi_i|$ and $\bar{\phi}_i$ for the six attributes characterizing a prospective LP, as reported in Fig. 3. For the GNN predictor, we also include the aggregate $|\phi_i|$ and $\bar{\phi}_i$ computed over the set of features characterizing the network topology embeddings and the slots occupation. Feature nomenclature is detailed in Tab. III. We now analyze the most important features.

1) *Path Length*: As expected, path length is the dominant feature for all scenarios and both models, as amplified spontaneous emission and nonlinear interference accumulate with distance [18]. However, for XGB, path length dominates with a SHAP range spanning ± 10 dB, followed by carrier frequency and length of the longest link spanning ± 2 -4 dB, while all other features cluster near zero. The GNN distributes importance more evenly across LP slot features, frequency, length of longest link, and number of links, as shown in Fig. 3.

2) *Band-specific feature attribution*: In single-band scenarios, carrier frequency exhibits a moderate monotonic effect reflecting first-fit allocation: higher-frequency slots are less loaded and thus experience higher SNR. In multiband operation, the mean $|\phi|$ of the carrier frequency feature markedly increases for GNN (Fig. 3, Multiband column), acting as discriminator between C-band and L-band operating conditions. More in detail, higher-frequency (C-band) channels contribute negatively to SNR, while lower-frequency (L-band) channels contribute positively. This reflects the cumulative effect of cross-band stimulated Raman scattering [1], [19], which transfers power from C to L band. However, channels at the lower-frequency L-band edge benefit more, due to cumulative Raman gain from both bands, whereas those near the C-L boundary receive weaker amplification. In the Ensemble scenario, the carrier frequency attribution loses its clean trend due to the superposition of two single-band decisional logics, whereas XGB maintains an unchanged pattern in all scenarios. This indicates that the GNN has learned a physically meaningful




Fig. 2. SHAP beeswarm plots for GNN (top) and XGB (bottom) for the four evaluation scenarios: A→B and AB→C (C-band only), ABC→D (multiband C+L) in green, and ABC/C→D (multiband C,L) via Ensemble model (one per band) in gold. Features are ranked by $\text{mean}|\phi_i|$; dot color encodes feature value.




Fig. 3. Mean absolute SHAP values $|\phi_i|$ per feature and scenario for GNN (top) and XGB (bottom).

band discriminator. Its sharp SHAP transition at the C/L boundary mirrors the Raman power-transfer gradient, while XGB treats frequency merely as an intra-band ordering cue, missing the inter-band mechanism entirely.

3) *LP Slots and Topology-Aware Embeddings (GNN)*: As reported in Fig. 3, LP slot descriptors encoding per-

link spectral occupation carry the highest attribution in the GNN model, even surpassing the path length attribute. Unlike XGB, which uses only attributes characterizing the spectrally-closest neighbors, the GNN processes the full spectral state of every traversed link, confirming that global spectral context is structurally integrated in the model. The topology-aware embeddings likewise show consistent SHAP contributions, ranking third (after LP slots and path length attributes) and thus providing empirical evidence that message passing captures topology-induced relational dependencies.

Notably, a consistent core of bottleneck links (Links # 5, 8, 9, 19, 24, 25) appears as high-importance features for all GNN scenarios. These links are traversed by a large share

TABLE III
FEATURE NOMENCLATURE USED IN THE SHAP ANALYSIS

Feature	Description
Path length	Total physical length of the LP route (km)
Num links	Number of fiber spans traversed by the LP
Longest link	Length of the longest individual span (km)
Frequency	Central carrier frequency of the LP (THz)
Traffic	Requested data rate (Gbps)
Tx Mode	Modulation format index
Link k	Per-slot spectral-occupation descriptor of link k , of the occupying LP (zero if empty)
LP slots	Aggregate $ \phi_i $ over all Link k – Slot features
GNN output	Path embedding from mean-pooling the GCN hidden states over traversed links
Neighbor features	Aggregate $ \phi_i $ over the 14 left/right spectral-neighbor descriptors




Fig. 4. Mean signed SHAP values $\bar{\phi}_i$ for GNN (left) and XGBoost (right) across all scenarios, showing both magnitude and direction of feature attributions.

of shortest-path routes, making them congestion points [19], [20]. Their stability across all training/test splits confirms the GNN is exploiting the network's structural load distribution rather than memorizing scenario-specific patterns, enabling operators to directly identify which links drive SNR degradation. This is confirmed by the aggregate $|\bar{\phi}_i|$ computed over all link slot features, which overcome the path length value in all scenarios (see Fig. 3, top). This contrast reveals a fundamental architectural difference: because XGB must compress all distance-related impairments into a single scalar, path length absorbs variance that physically originates from per-link spectral loading and gain tilt; the GNN can decouple it thanks to the per-link spectral state info.

Furthermore, the GNN beeswarms reveal band-specific link hotspots beyond the universal bottleneck core: links $\{1, 2, 6, 21, 31, 42\}$ appear selectively in C-band scenarios (Fig. 5, blue links), while links $\{7, 10, 12, 13, 40, 41\}$ emerge exclusively in the L-band model (Fig. 5, yellow links), as visible in Fig. 2. This band-selective activation demonstrates that the GNN learns not only *which* links matter but *why* they matter differently across bands, providing operators with actionable diagnostics: universal bottlenecks (red) flag links requiring capacity upgrades regardless of band, whereas L-band hotspots (yellow) pinpoint where cross-band Raman mitigation [19] would yield the highest benefit.

D. Feature Attribution Stability

Tab. IV reports both the concentration ratio C and the Spearman rank correlation ρ [15], [16], which measures the rank correlation between the importance vectors $|\bar{\phi}_i|_i$ of consecutive scenarios (e.g., AB→C vs. A→B). For the L-band, ρ is computed against the last C-band scenario (ABC→D).

XGB exhibits $C \approx 0.60$ in all scenarios with $\rho = 1.0$ between consecutive C-band pairs, indicating that both the concentration and the ranking of feature attributions remain unchanged as load increases. This suggests that its explanatory




Fig. 5. Physical network topology with directed link IDs. Link colors indicate SHAP-identified high-importance links: universal bottlenecks (red), C-band-specific (blue) and L-band-specific hotspots (yellow).

logic is largely insensitive to evolving spectral conditions. In contrast, the GNN shows lower and more variable concentration ($C \approx 0.52\text{--}0.56$) together with reduced correlation across stages, indicating that it adapts its attribution structure as the network transitions through different traffic loads.

The inter-band comparison is the most operationally significant result. When moving from the final C-band scenario to the L-band model, the GNN yields $\rho = 0.29$, indicating

TABLE IV

ATTRIBUTION STABILITY METRICS. C = CONCENTRATION RATIO (EQ. 1);
 ρ = SPEARMAN CORRELATION WITH PRECEDING C-BAND SCENARIO;
 † C-VS-L BAND-SPECIFIC COMPARISON.

Model	Scenario	C	ρ
GNN	A→B (C-band)	0.52	–
	AB→C (C-band)	0.56	0.95
	ABC→D (C-band)	0.54	0.71
	C→D (L-band)	0.52	0.29[†]
XGB	A→B (C-band)	0.60	–
	AB→C (C-band)	0.53	1.00
	ABC→D (C-band)	0.60	1.00
	C→D (L-band)	0.61	0.94[†]

a substantial reordering of feature importance and thus a band-specific explanatory structure. By contrast, XGB yields $\rho = 0.94$, meaning that its feature rankings remain nearly unchanged across the band transition, with a maximum rank displacement of 1 versus a displacement of 6 for the number of links attribute in the GNN. This near invariance suggests that XGB essentially applies the same decision logic in both bands, providing false confidence that identical impairment mechanisms govern C- and L-band behavior, whereas the GNN captures the expected shift in underlying physical effects. Fig. 4 substantiates these findings: XGB maintains nearly identical signed contributions across scenarios (except for sign inversions of carrier frequency in the multiband model and path length in the Ensemble model), whereas the GNN exhibits substantial shifts in both magnitude and sign. This rigidity of XGB’s attribution structure provides the explainability-level evidence for the systematic long-haul overestimation reported in [4]: by applying C-band logic unchanged to L-band lightpaths, XGB cannot account for the changes in the impairments’ impact, whereas the GNN’s adaptive feature hierarchy avoids such downside. This finding is operationally significant: as discussed in [4], SNR overestimation risks service degradation, whereas the GNN’s more conservative estimation provides a safety margin for long-haul provisioning.

V. CONCLUSIONS

This paper presented a SHAP-based explainability comparison between a topology-aware GNN and a feature-based XGB model for QoT estimation in multiband optical networks. XGB concentrates $\approx 60\%$ of its attribution budget on path length across scenarios, a decisional logic that does not adapt to traffic load changes or detect band-specific impairment mechanisms. Conversely, the GNN distributes importance across topology-aware embeddings, per-link spectral occupancy, and routing descriptors, identifies recurring congestion bottlenecks in the physical topology, and better captures cross-band Raman interactions at the C/L band boundary.

REFERENCES

- [1] N. Deng, L. Zong, H. Jiang, Y. Duan, and K. Zhang, “Challenges and enabling technologies for multi-band WDM optical networks,” *Journal of Lightwave Technology*, vol. 40, no. 11, pp. 3385–3394, 2022.
- [2] F. Musumeci, C. Rottondi, A. Nag, I. Macaluso, D. Zibar, M. Ruffini, and M. Tornatore, “An overview on application of machine learning techniques in optical networks,” *IEEE Communications Surveys & Tutorials*, vol. 21, no. 2, pp. 1383–1408, 2019.
- [3] R. Ayassi, A. Triki, N. Crespi, R. Minerva, and M. Laye, “Survey on the use of machine learning for quality of transmission estimation in optical transport networks,” *Journal of Lightwave Technology*, vol. 40, no. 17, pp. 5803–5815, 2022.
- [4] J. Eid, F. Galante, O. Ayoub, G. S. Sticca, M. Ibrahimi, A. Bianco, M. Tornatore, and C. Rottondi, “GNN vs. XGBoost for QoT Estimation in Incremental C+L Multiband Networks,” in *Proceedings of the International Conference on Optical Network Design and Modelling (ONDM)*, 2026.
- [5] Y. Pointurier, “Machine learning techniques for quality of transmission estimation in optical networks,” *Journal of Optical Communications and Networking*, vol. 13, no. 4, pp. B60–B71, Apr. 2021.
- [6] O. Ayoub, S. Troia, D. Andreoletti, A. Bianco, M. Tornatore, S. Giordano, and C. Rottondi, “Towards explainable artificial intelligence in optical networks: The use case of lightpath QoT estimation,” *Journal of Optical Communications and Networking*, vol. 15, no. 1, pp. A26–A38, Jan. 2023.
- [7] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, “A comprehensive survey on graph neural networks,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [8] T. Panayiotou, G. Savva, B. Shariati, I. Tomkos, and G. Ellinas, “Machine learning for QoT estimation of unseen optical network states,” in *Optical Fiber Communications Conference and Exhibition (OFC)*, 2019, pp. 1–3.
- [9] G. Savva, T. Panayiotou, I. Tomkos, and G. Ellinas, “Deep graph learning for QoT estimation of unseen optical sub-network states: Capturing the crosstalk impact on the in-service lightpaths,” *Journal of Lightwave Technology*, vol. 40, no. 4, pp. 921–934, Feb. 2022.
- [10] A. Prakash and S. Kar, “Supervised graph convolution networks for OSNR and power estimation in optical mesh networks,” *Journal of Optical Communications and Networking*, vol. 14, no. 6, pp. 469–480, 2022.
- [11] H. Houssiany, O. Ayoub, C. Rottondi, and A. Bianco, “Using shap values to validate model’s uncertain decision for ml-based lightpath quality-of-transmission estimation,” in *2023 23rd International Conference on Transparent Optical Networks (ICTON)*, 2023, pp. 1–5.
- [12] A. Fawaz, F. Arpanaei, H. Beyranvand, C. Natalino, J. M. Rivas-Moscoco, D. Larrabeiti, A. Napoli, and J. P. Fernandez-Palacios, “Reducing complexity and enhancing predictive power of ML-based lightpath QoT estimation via SHAP-assisted feature selection,” in *Intl. Conf. on Optical Network Design and Modelling (ONDM)*, 2024.
- [13] Z. Zhang, X. Tian, T. Ma, Z. Zheng, and J. Hu, “SHAP-assisted EE-LightGBM model for explainable fault diagnosis in practical optical networks,” *Journal of Optical Communications and Networking*, vol. 17, no. 2, pp. 81–94, Feb. 2025.
- [14] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” *Advances in neural information processing systems*, vol. 30, 2017.
- [15] C. Spearman, *The Proof and Measurement of Association between Two Things*, 1904, vol. 15.
- [16] J. Yuan and A. Dasgupta, “Fooling SHAP with output shuffling attacks,” in *AAAI*, 2024.
- [17] M. Ibrahimi, H. Abdollahi, C. Rottondi, A. Giusti, A. Ferrari, V. Curri, and M. Tornatore, “Machine learning regression for QoT estimation of unestablished lightpaths,” *Journal of Optical Communications and Networking*, vol. 13, no. 4, pp. B92–B101, 2021.
- [18] D. Semrau, R. I. Killey, and P. Bayvel, “A Closed-Form Approximation of the Gaussian Noise Model in the Presence of Inter-Channel Stimulated Raman Scattering,” *IEEE Journal of Lightwave Technology*, vol. 37, no. 9, pp. 1924–1936, 2019.
- [19] G. Sticca, M. Ibrahimi, F. Musumeci, N. Di Cicco, A. Castoldi, R. Pastorelli, and M. Tornatore, “Selective hybrid EDFA/Raman amplifier placement to mitigate lightpath degradation in (C+L) networks,” *Journal of Optical Communications and Networking*, vol. 15, no. 8, pp. C232–C241, 2023.
- [20] G. S. Sticca, M. Ibrahimi, N. Di Cicco, F. Musumeci, and M. Tornatore, “Incremental planning with dual-fiber distributed Raman amplification in (C+L+S) networks,” *Journal of Optical Communications and Networking*, vol. 17, no. 9, pp. D156–D166, 2025.