



## Policy-driven Conformal Prediction for Trustworthy QoT Estimation

Downloaded from: <https://research.chalmers.se>, 2026-07-12 18:11 UTC

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

Rezaei, K., Ayoub, O., Monti, P. et al (2026). Policy-driven Conformal Prediction for Trustworthy QoT Estimation. Proceedings of the Optical Fiber Communication Conference (OFC) 2026.  
<http://dx.doi.org/10.1364/OFC.2026.M4A.3>

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

# Policy-driven Conformal Prediction for Trustworthy QoT Estimation

Kiarash Rezaei <sup>1,\*</sup> Omran Ayoub <sup>2</sup> Paolo Monti <sup>1</sup> and Carlos Natalino <sup>1</sup>

<sup>1</sup>Department of Electrical Engineering, Chalmers University of Technology, 412 96 Gothenburg, Sweden

<sup>2</sup>University of Applied Sciences and Arts of Southern Switzerland, 6928 Lugano, Switzerland

\*kiarashr@chalmers.se

**Abstract:** We propose Conformal QoT, a policy-driven framework that combines statistically guaranteed QoT estimation with operational decision policies, enabling reliable lightpath-feasibility predictions under domain shift and improving accuracy from 92% to 99.6% on open datasets.

## 1. Introduction

Trustworthy quality of transmission (QoT) estimation is essential for informed optical network operations, particularly as networks grow more complex through disaggregation, spatial multiplexing (multi-core, multi-band), and tighter design margins. In this evolving landscape, trustworthiness implies not only accuracy but also consistent and reliable predictions across domains and under varying network conditions, often with scarce available data, as improved estimation directly translates to reduced costs and extended network lifetime. One of the most important aspects of trustworthy artificial intelligence & machine learning (AI/ML) is *uncertainty quantification*, which enables assessing the confidence of its predictions and identify cases where decisions should be treated with caution. A model should also adapt to changing conditions without requiring extensive data. Thus, a model that indicates when and where its predictions may be unreliable is inherently more trustworthy. Yet, most existing machine learning (ML)-based models do not effectively quantify uncertainty, and this lack of quantification undermines the trust and adoption of data-driven optical network design, planning, management, and control [1].

Early studies framed the ML-based QoT estimation task as binary feasibility classification [2, 3], later extending to uncertainty-aware approaches where neural networks were used to jointly estimate optical signal-to-noise ratio (OSNR) and its uncertainty via dropout, identifying unreliable predictions [4]. Bayesian and quantile regression further advanced probabilistic QoT forecasting [5]. Calibrated regression with gradient boosting and isotonic recalibration has produced statistically consistent intervals [6], and probabilistic low-margin design [7] has balanced spectrum use and risk. However, these approaches remain model- and data-dependent, limiting their generalizability. As a result, their performance often degrades under changing network conditions or across different network domains, limiting practical use for operators who have limited data and need QoT estimates without frequent retraining or recalibration.

The paper addresses the above challenges by proposing a two-layer, policy-driven framework, referred to as *Conformal QoT*, which combines statistically guaranteed QoT estimation with its translation into operational decisions. In the first layer, conformal prediction (CP) is applied to the ML-based QoT estimation, providing model-agnostic, distribution-free guarantees that remain valid under domain shift without retraining. Each QoT prediction is accompanied by a confidence band with formal guarantees, enabling trustworthy assessment of prediction reliability. The second layer introduces a decision policy that combines the QoT estimation and conformal confidence bands into reliable lightpath (LP) feasibility decisions, linking predictive confidence with operational action. We assess the framework in realistic deployment scenarios where models trained on data of one transmission technology are reused with a different one. Results show that combining guaranteed QoT estimation with our policy-driven decision-making improves accuracy from 92% to 99.6% on the same dataset, enabling tighter margins that reduce infrastructure costs and extend network lifetime. For cross-domain deployment, a model trained on source-domain data and conformalized with only 20% of target-domain samples achieves a 21.7% accuracy improvement over the baseline while reducing false positives by over 99%. This demonstrates that the framework enables reliable model reuse across network domains with minimal target data and without retraining, effectively bridging the gap between uncertainty estimation and operational decision-making in modern optical networks.

## 2. Conformal QoT Framework

The proposed *Conformal QoT* framework (Fig. 1) comprises three parts: (i) a ML-based QoT regressor; (ii) a CP that produces statistically reliable confidence bands; and (iii) a policy that drives the LP feasibility decisions. Operations begin with a base AI/ML regressor that predicts the QoT metric of interest for a given LP sample (Fig. 1(a)). Traditionally, a model uncertainty (MU) band (Fig. 1(b)) can be extracted from a AI/ML regressor, quantifying how predictions vary when the model is run multiple times with different random seeds, indicating its confidence level. However, the MU band lacks statistical assurance, as the true QoT value may fall outside its




Fig. 1. Overview of the Conformal QoT framework.

range. Moreover, the MU band is dataset-dependent, making it unreliable under domain shift, i.e., when network conditions change.

The *CP layer* (Fig. 1(c)) overcomes these limitations by providing distribution-free statistical guarantees that remain valid even under domain shift. It uses a small conformity set from the target domain validation data to assess deviations between predicted and true QoT values. This approach remains effective even with limited target-domain data, as a small conformity set still provides formal guarantees. Using these deviations and a user-defined significance level  $\alpha$  (e.g.,  $\alpha = 0.1$  for 90% confidence), we compute an error bound such that at least  $(1 - \alpha)$  of the samples lie within the bound. This bound extends each prediction to form the global conformal (GC) band (Fig. 1(d)), ensuring that the true QoT lies within it with probability  $(1 - \alpha)$ , without distributional assumptions or retraining. The *policy layer* (Fig. 1(e)) translates the statistically valid GC bands into LP feasibility decisions through a tunable risk-tolerance parameter  $\delta$ , enabling operators to balance spectrum efficiency against operational reliability based on their specific requirements. We propose two policies that leverage the GC bands differently. The risk-averse (RA) policy considers a LP feasible only if the entire GC band lies above the QoT feasibility threshold, prioritizing conservative decisions that minimize incorrect acceptances (false positives). In contrast, the efficiency-oriented (EO) policy classifies a LP as feasible when the upper bound of the band exceeds the threshold plus  $\delta$ , enabling more aggressive provisioning while maintaining statistical guarantees.

### 3. Experiments and Results

We evaluated the *Conformal QoT* framework using public datasets from [8] to demonstrate its effectiveness under domain shift. Dataset 01 (Bu-SMF) and Dataset 07 (WC-MCF) from the CONUS topology represent source and target domains for single-core and multi-core fiber technologies, emulating scenarios where operators deploy new transmission technologies and assess whether models trained on the existing infrastructure can be reliably reused. We adopt an XGBoost (XGB) model as the base regressor due to its strong performance in prior QoT studies, training it to predict the received OSNR of each LP configuration. Point predictions serve as the baseline policy, where decisions rely solely on predicted OSNR values and feasibility thresholds without uncertainty quantification.

Three training regimes were evaluated to assess cross-domain generalization capabilities: zero-shot (ZS), transfer learning (TL), and target-domain (TD). In ZS, the model is trained on 80% of source domain samples (single-core) and evaluated directly on the target domain (multi-core), with 20% of target data reserved for the conformity set, simulating operators reusing existing models on new infrastructure without retraining. For TL, the model is pre-trained on the source domain, then fine-tuned using 20% of target data. In TD, the model is trained and tested entirely within the target domain using 90% for training and 10% for the conformity set. All data splits were mutually exclusive to prevent leakage.

Table 1. Regression (mean absolute error (MAE) and root mean squared error (RMSE)) and uncertainty (model uncertainty (MU) and global conformal (GC)) metrics across training regimes.

| Model                  | Regression Metrics |           | MU        |            | GC        |            |
|------------------------|--------------------|-----------|-----------|------------|-----------|------------|
|                        | MAE [dB]           | RMSE [dB] | Conf. [%] | Width [dB] | Conf. [%] | Width [dB] |
| Zero-shot (ZS)         | 1.97               | 1.98      | 0         | 0.04       | 90        | 4.17       |
| Transfer learning (TL) | 0.03               | 0.03      | 18        | 0.01       | 90        | 0.11       |
| Target-domain (TD)     | 0.02               | 0.03      | 32        | 0.02       | 91        | 0.10       |

Table 1 shows that while XGB achieves excellent accuracy in TL and TD regimes (sub-dB MAE), the ZS model shows significant degradation due to domain mismatch. The MU bands prove unreliable across all regimes, achieving only 0%, 18%, and 32% confidence despite narrow widths that misleadingly suggest certainty. This demonstrates why traditional uncertainty estimates fail under domain shift and cannot be trusted for operational decisions. In contrast, conformalized GC bands consistently deliver 90–91% confidence across all regimes. For TL and TD, these guarantees come with minimal band widths, enabling tight margins that translate into reliable LP decisions. Even in the challenging ZS scenario, the framework provides valid guarantees with wider bands (4.17 dB), reflecting inherent cross-domain uncertainty.

To assess the reliability of LP feasibility decisions, each test LP was evaluated with its assigned modulation format and associated OSNR threshold. Table 2 presents detailed classification metrics, while Fig. 2 summarizes

Table 2. Performance results and risk-tolerance parameters  $\delta$ .

| Model                  | Point + Baseline |      |      |     | GC + EO( $\delta_{EO}$ ) |      |     |     | GC + RA( $\delta_{RA}$ ) |      |     |      | $\delta_{EO}$ | $\delta_{RA}$ |
|------------------------|------------------|------|------|-----|--------------------------|------|-----|-----|--------------------------|------|-----|------|---------------|---------------|
|                        | TPR              | TNR  | FPR  | FNR | TPR                      | TNR  | FPR | FNR | TPR                      | TNR  | FPR | FNR  |               |               |
| Zero-shot (ZS)         | 29.4             | 48.3 | 22.3 | 0.0 | 62.2                     | 37.2 | 0.3 | 0.3 | 32.9                     | 37.5 | 0.0 | 29.7 | 0.1           | 0.0           |
| Transfer learning (TL) | 29.4             | 62.5 | 8.0  | 0.0 | 62.1                     | 37.5 | 0.0 | 0.4 | 62.0                     | 37.5 | 0.0 | 0.5  | 0.1           | 0.0           |
| Target-domain (TD)     | 29.4             | 62.6 | 8.0  | 0.0 | 62.5                     | 37.1 | 0.4 | 0.0 | 62.1                     | 37.5 | 0.0 | 0.4  | 0.0           | 0.0           |




Fig. 2. Accuracy (Acc), recall (Re), and F1-score (F1) across training regimes.

overall performance. True positive rate (TPR) and true negative rate (TNR) indicate correctly feasible and correctly infeasible LPs, while false positive rate (FPR) and false negative rate (FNR) capture incorrect decisions leading to either failed lightpath establishments (operational disruptions) or missed provisioning opportunities (lower spectral efficiency or revenue loss).

In Table 2, the operational impact is most evident in ZS, where baseline predictions yield high FPR, i.e., nearly one in four feasible lightpaths would fail. The EO policy reduces this failure rate by over 99% (to 0.3%) while more than doubling TPR (from 29.4% to 62.2%). The RA policy eliminates false acceptances (0% FPR) with moderate conservatism (FNR  $\approx$  29.7%), offering maximum operational safety. This demonstrates that models trained on existing infrastructure can be reliably deployed on new technologies using only a small conformity set, eliminating the need for retraining and enabling rapid technology adoption. For TL and TD models, both policies maintain high TPR (around 62%) with near-zero FNR while reducing FPR from 8% to below 0.5%. This transforms nominally accurate models (92% baseline accuracy) into highly reliable systems (99.6% accuracy), enabling tighter margins that extend network lifetime and reduce infrastructure costs. The improvement from 77.7% to 99.4% accuracy in cross-domain ZS is particularly significant, confirming the framework enables model reuse across network domains with limited target data and without retraining. Fig. 2 shows that conformalized policies consistently improve F1-score, recall, and accuracy across all regimes. Notably, the conformalized ZS model with EO policy achieves performance comparable to TL and TD models, demonstrating that statistical guarantees through conformal prediction bridge the domain gap and enable reliable zero-shot deployment, critical as optical networks evolve toward disaggregated, multi-band, multi-core architectures where training data is scarce. Operators can adjust  $(1 - \alpha)$  and  $\delta$  to balance reliability against efficiency based on their operational requirements.

#### 4. Conclusions

We combine CP with policy-driven decision logic to provide statistically guaranteed confidence bands that remain valid under domain shift without retraining. It reduces LP decision errors and enables reliable cross-domain model reuse. By delivering trustworthy QoT estimation with minimal target-domain data, the framework supports rapid and reliable technology deployment in modern optical networks.

**Acknowledgments:** This work has been supported by the Horizon Europe ECO-eNET project with funding from the SNS JU under grant agreement No. 101139133.

#### References

1. F. Musumeci, *et al.*, J. Opt. Commun. Netw. **17**, C144–C155 (2025). DOI: [10.1364/JOCN.551910](https://doi.org/10.1364/JOCN.551910).
2. T. Panayiotou, *et al.*, J. Opt. Commun. Netw. **9**, 98–108 (2017). DOI: [10.1364/JOCN.9.000098](https://doi.org/10.1364/JOCN.9.000098).
3. C. Rottondi, *et al.*, J. Opt. Commun. Netw. **10**, A286–A297 (2018). DOI: [10.1364/JOCN.10.00A286](https://doi.org/10.1364/JOCN.10.00A286).
4. T. Tanimura, *et al.*, J. Light. Technol. **37**, 1717–1723 (2019). DOI: [10.1109/JLT.2019.2895730](https://doi.org/10.1109/JLT.2019.2895730).
5. F. Meng, *et al.*, J. Opt. Commun. Netw. **11**, A144–A154 (2019). DOI: [10.1364/JOCN.11.00A144](https://doi.org/10.1364/JOCN.11.00A144).
6. N. Di Cicco, *et al.*, in *Proc. of BalkanCom*, (2022), pp. 21–25. DOI: [10.1109/BalkanCom55633.2022.9900791](https://doi.org/10.1109/BalkanCom55633.2022.9900791).
7. O. Karandin, *et al.*, J. Opt. Commun. Netw. **15**, C129–C137 (2023). DOI: [10.1364/JOCN.482734](https://doi.org/10.1364/JOCN.482734).
8. H. Akbari, *et al.*, J. Opt. Commun. Netw. **17**, 514–525 (2025). DOI: [10.1364/JOCN.558452](https://doi.org/10.1364/JOCN.558452).