Computer Networks 212 (2022) 108992

Contents lists available at ScienceDirect

Computer Networks

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

Learning quantile QoT models to address uncertainty over unseen lightpaths

Hafsa Maryam, Tania Panayiotou *, Georgios Ellinas

Department of Electrical and Computer Engineering and KIOS Research and Innovation Center of Excellence, University of Cyprus, Cyprus

ARTICLE INFO

Keywords:
QoT estimation
Quantile regression
Machine learning
Margin reduction
Resource allocation
Optical networks

ABSTRACT

Uncertainty in quality-of-transmission (QoT) estimation is traditionally addressed through empirical, myopic margins, ignoring the fact that each unseen lightpath is subject to different levels of uncertainty. To address this limitation, in this work, deep quantile regression is leveraged to finer capture QoT estimation uncertainty through the inference of margins that act discriminative over the unseen lightpaths. Specifically, deep-quantile regression is applied to approximate QoT models capable of inferring the QoT of unseen lightpaths, according to a predefined level of certainty. Quantile models automatically account for the uncertainty during inference, without the need to consider additional empirical margins for decision-making (i.e., the margins are learned and considered upon inference). It is shown that quantile QoT models lead to significant margin reduction when compared to baseline myopic margin schemes, resulting in more confident and network efficient allocation decisions.

1. Introduction

Machine learning (ML), with proven capabilities on sufficiently modeling the non-linear nature of physical layer impairments (PLIs), has been extensively studied in the optical networks literature with the purpose of accurately estimating the quality-of-transmission (QoT) of unseen (unestablished) lightpaths [1–4]. This information, is then used to examine the feasibility of any unseen lightpath and take decisions accordingly; that is, establish or deny service to the unseen lightpath under consideration.

Commonly, regression is leveraged to approximate a QoT model capable of conducting inference about the QoT value of unseen lightpaths. Specifically, with regression, a least squares loss function is usually minimized to find a QoT model, which is, in essence, an approximation of the relationship between the input variables (i.e., lightpaths) and the output variables (i.e., QoT values). The output of the QoT model is an estimate or an approximation containing some uncertainty that is represented by the variability around the mean response value. Obviously, model accuracy, even though important, is not a sufficient measure of how much such point-based estimates can be trusted for confident decision-making. In fact, uncertainty, resulting from the errors in the model itself and the noise over the input variables, may result in erroneous decisions, especially when the QoT estimates greatly overestimate the true QoT values of some lightpaths; that is, a lightpath with a true QoT value that violates some acceptable threshold may be erroneously considered feasible and hence admitted in the network instead of being blocked.

While QoT model uncertainty is already considered in the literature, this is generally done through the estimation of myopic empirical margins [5,6], largely ignoring the fact that diverse input patterns are subject to different levels of uncertainty. One limitation of considering such margins, is that the QoT estimate of every unseen lightpath is equally degraded, without discrimination, according to a value that reflects the worst uncertainty level that can be statistically measured. While by doing so it is ensured that lightpaths with a true QoT that is insufficient will not be established in the network (i.e., the possibility of establishing in the network a lightpath with insufficient QoT is mitigated or even alleviated), the true QoT of some feasible lightpaths may be greatly underestimated, possibly leading to unjustified denial of service, lightpath overprovisioning, and ultimately to inefficient network decisions and waste of resources.

To this end, given the existing gap in the literature on closely capturing QoT estimation uncertainty (i.e., fine-tuning margins over lightpath uncertainty), the main contribution of this work is to provide a framework capable of measuring the certainty level of QoT model estimates towards confident and network efficient decision-making. In ML, several approaches exist capable of achieving this objective, including conducting Bayesian or Monte Carlo dropout inference [7], and through the estimation of conditional quantile functions (e.g., deep quantile regression) [8]. In this work, as a first step towards capturing QoT estimation uncertainty, a deep quantile regression framework is adopted, in which margins are learned and inferred in a discriminative fashion over lightpath patterns. In this framework, deep neural

* Corresponding author.
E-mail addresses: maryam.hafsa@ucy.ac.cy (H. Maryam), panayiotou.tania@ucy.ac.cy (T. Panayiotou), gellinas@ucy.ac.cy (G. Ellinas).

https://doi.org/10.1016/j.comnet.2022.108992
Received 18 January 2022; Received in revised form 29 March 2022; Accepted 19 April 2022
Available online 27 April 2022

networks (DNNs) are trained to estimate conditional quantile functions (i.e., models) by minimizing an asymmetrically weighted sum of absolute errors [9]. Hence, the outputs of estimated conditional quantile functions automatically take the uncertainty into consideration and can be directly used for decision-making without considering any additional margins.

1.1. Related work

Machine learning has been widely adopted in the optical networks literature [10,11] and within the communications systems in general [12,13] for various functions and applications. Regarding optical networks, where this work focuses on, ML has been indicatively applied for traffic prediction and resource allocation optimization [14–16], fault detection/localization [17–19], attack detection/identification [20,21], and QoT estimation [1,5,6,22]. In this work, we focus on ML for QoT estimation for which a survey of the state-of-the-art can be found in [23]. Given the already rich literature on ML for QoT estimation, evidently, this topic has received considerable attention over the last few years.

In fact, ML is considered today as a promising approach for modeling the non-linear nature of PLIs present in optical networks (e.g., crosstalk, polarization mode dispersion, amplified spontaneous emission noise, etc.), demonstrating significant margin reduction compared to the conservative margins considered in traditional linear physical layer models (PLMs) [24–26]. A general ML-aided QoT estimation framework, enabled by the software-defined networking (SDN) advances [15,27], can be described as the exploitation of historical QoT data of previously established lightpaths with the purpose of inferring accurate QoT models for unseen (unestablished) lightpaths [2,3,6,22,28–30].

To achieve this objective, several ML models have been exploited, such as DNNs, k-nearest neighbors, support vector machines, and Gaussian Processes [23,31], with DNNs validated both in the field [32,33] and experimentally [2,34] with the use of real datasets. In general, the model of choice, is trained either as a binary classifier [29,35–37] or as a regressor [5,6,22,27] by means of supervised learning, given a labeled dataset of previously observed lightpaths.

In classification, the model is trained to predict the class of an unseen lightpath; that is, the feasible or the infeasible class, which is determined merely according to a predefined QoT threshold. In essence, classification combines prediction and decision-making, which, may however lead to erroneous decisions. Specifically, while it was shown that high classification accuracies can be achieved in both classes of interest, the slight model inaccuracies (i.e., uncertainty) that still exist within each class may lead to erroneous decisions, especially when an infeasible lightpath is misclassified to the class of feasible lightpaths.

As classification does not provide any way of fixing such inaccuracies prior to decision-making, regression is applied to infer the actual QoT values (i.e., bit error rate, optical signal to noise ration, Q-factor) of unseen lightpaths [5,6,22,27]. Regression, enables the estimation of error margins, that when considered before decision-making can, to some extend, fix model inaccuracies. Commonly, such margins are used to degrade the inferred QoT values of every unseen lightpath prior to deciding about their feasibility.

Related works addressing model inaccuracies with error margins approximate a QoT model by minimizing a mean squared error (MSE) loss function over a training dataset. The test dataset is afterwards used to evaluate both the performance accuracy and error spread of the model [5,6], with the error spread used for margin estimation. However, as previously mentioned, in this approach a constant margin is estimated and considered, without discrimination, for all unseen lightpaths prior to decision-making. This approach may lead to erroneous decisions, especially for those unseen lightpaths that the model is almost certain about their feasibility, as diverse unseen lightpaths are subject to different levels of uncertainty.

1.2. Contribution

Evidently, QoT model margins are currently roughly approximated, with the potential of ML to achieve further improvement in decision-making and network efficiency remaining greatly unexplored. Thus, in order to appropriately represent and consider the QoT model and input uncertainty over unseen lightpaths, in this work, we present an uncertainty representation framework aiming to alleviate the limitations of existing myopic margins that are commonly applied to compensate for QoT model uncertainty.

Specifically, the aim is to closely capture model uncertainty by fine-tuning the estimation of margins that act discriminatively over the unseen lightpaths. In particular, we leverage deep quantile regression [8] to approximate conditional quantile functions capable of inferring QoT values that already account for margins, according to a predefined certainty level q; that is, given an unseen lightpath, the trained model infers a QoT value such that the probability that the true QoT of this lightpath is below the inferred value is equal to q. Hence, the inferred QoT values of each unseen lightpath are in essence automatically adjusted by the conditional quantile model (i.e., the q-quantile) to consider what the model already knows about each unseen lightpath.

Preliminary results of this work were first presented in [38], where it was shown that q-quantiles achieve significant margin reduction when compared to baseline margin estimation schemes without, however, any compromise on the classification accuracy achieved for both the classes of interest. In contrast, it was shown that baseline, myopic, margins greatly underestimate the true QoT of the lightpaths, leading to accurate decisions only in the class of infeasible lightpaths. These outcomes were obtained for a wavelength division multiplexing (WDM) network examining only a single lower q threshold. The QoT was measured in dB according to the Q-factor metric for WDM networks [26]. This work, greatly extends [38] by:

- Examining margin reduction and accuracy of decisions for various q-quantiles acting as lower estimation bounds. The purpose is to investigate the importance of appropriately fine-tuning the lower q value towards further margin reduction and higher classification accuracy.
- Examining the impact of the best q-quantile obtained on the performance of the network. Specifically, the performance of various margin schemes are evaluated and compared in terms of blocking probability on a dynamic optical network. In this network, an elastic optical network (EON) is considered, and routing and spectrum allocation (RSA) decisions are taken by utilizing: (i) the q-quantile model, (ii) the least-squares model considering error spread margin.
- Creating appropriate datasets for model training and evaluation, by redefining the input lightpath features of the models to now account for the EON technology. The QoT values are measured in dB according to the Q-factor metric for EONs [24].

It is shown that all q-quantiles considered achieve significant margin reduction compared to the baseline approaches, also demonstrating significant improvement in the overall classification accuracy and accuracy achieved for the class of feasible lightpaths. A small penalty is observed in the classification accuracy of infeasible lightpaths (i.e., up to 0.3%), which is however negligible given the overall improvements observed. Furthermore, it is shown that different q-quantiles perform differently, indicating the importance of fine-tuning parameter q before considering a model for decision-making. Importantly, it is shown that q-quantiles have a positive impact on network performance, significantly outperforming baseline approaches in terms of blocking.

The rest of the paper is organized as follows. Section 2 presents background information on deep quantile regression, Section 3 discusses dataset generation, and Section 4 presents the model training and evaluation procedure. The design margin estimation results are

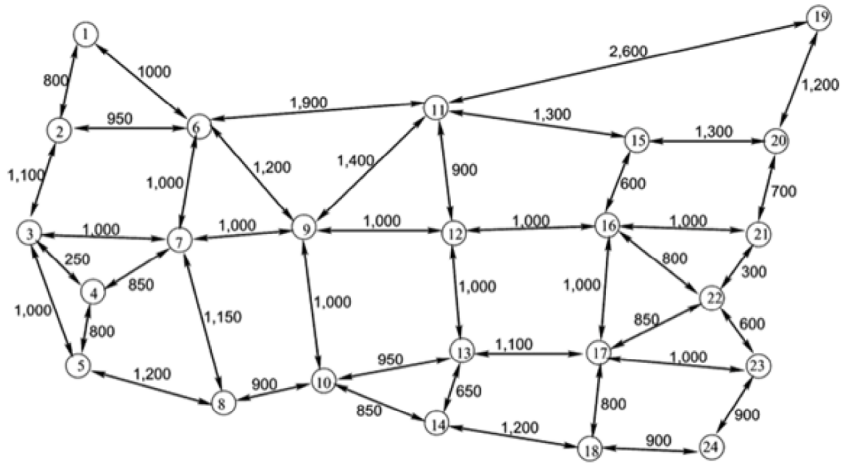

Fig. 1. USNET topology.

presented in Section 5, including a comparison with benchmark techniques, while the decision accuracy results and their comparison with the benchmark approaches are presented in Section 6. Finally, Section 7 examines the impact of the proposed approaches on the network performance in terms of blocking probability, and Section 8 presents concluding remarks and avenues for future research.

2. Deep quantile regression for QoT inference

In this work, deep quantile regression is applied to approximate the conditional quantile function $Q_Y(q|X)$, also known as the q-quantile, where $0 < q < 1$, X is a r.v. representing the past lightpath observations, Y is a r.v. representing the true QoT value of unseen lightpaths, and $\mathbf{x}$, y are possible realizations of X and Y. By definition, a q-quantile is learned to return a value $\hat{y}$, such that the probability that the true QoT value y of an unseen lightpath, given $\mathbf{x}$, will be less than or equal to $\hat{y}$, is equal to q [9]. In essence, the q-quantile determines the percentage of a population that is above or below a certain threshold. Hence, for our use case, where the QoT value is measured in dB, it is important that q is low enough (i.e., $q \to 0$) to approximate a lower QoT estimation bound $\hat{y}$ for each input $\mathbf{x}$; that is, to increase the probability that the true QoT of an unseen lightpath $\mathbf{x}$ does not fall below QoT estimate $\hat{y}$.

In a deep learning framework, a q-quantile is estimated by minimizing the asymmetrically weighted sum of absolute errors [8,9]:

$$L_q = \frac{1}{n}\sum_{i=1}^{n} \rho_q\big(y_i - \hat{Q}_Y(q|\mathbf{x}_i)\big), \quad \rho_q(z) = \begin{cases} qz, & \text{if } z \geq 0, \\ (q-1)z, & \text{if } z < 0, \end{cases} \tag{1}$$

where n is the number of observed lightpaths, $\mathbf{x}_i \in \mathbb{R}^d$ is a vector describing the ith lightpath of X, $y_i \in \mathbb{R}$ is the true QoT value of this lightpath, and $\hat{Q}_Y(q|\mathbf{x}_i)$ is an approximation of the q-quantile returning QoT estimate $\hat{y}_i$ for lightpath i, $\forall\, i = 1, \ldots, n$. In this work, the q-quantile is parameterized by a DNN model that is trained to estimate $\hat{Q}_Y(q|X, \theta)$, where θ are the unknown parameters of the quantile DNN model. Specifically, the model is optimized to minimize the loss function of Eq. (1) given a labeled dataset $D = (X, Y) = \{\mathbf{x}_i, y_i\}_{i=1}^{n}$.

In dataset D, the inputs of the DNN model are defined according to an EON over the feature vectors $\mathbf{x}_i = [x_{i,1}, \ldots, x_{i,7}]$, where $x_{i,1}$ is the length (in km) of lightpath i, $x_{i,2}$ is the maximum link length (in km) of i, $x_{i,3}$ is the identification number of the first spectrum slot allocated to i, $x_{i,4}$ is the number of spectrum slots allocated to i, $x_{i,5} = 1, 2, 3, 4$ if BPSK, QPSK, 8-QAM, or 16-QAM modulation format is used for i, respectively, $x_{i,6}$ is the number of erbium doped fiber amplifiers (EDFAs) along i, and $x_{i,7}$ is the number of links along i. Note that these inputs constitute commonly considered lightpath features that have been proven to sufficiently describe a lightpath for ML-aided QoT estimation [29,30].

Regarding the output of the DNN model, y_i, this is the Q-factor value of i that in practice can be obtained through optical performance monitoring (OPM) [15]. Since OPM can only take place after a lightpath is established, probing lightpaths can be used to enrich the dataset information with lightpaths of insufficient QoT values (i.e., lightpaths that cannot be admitted in the network). As probing lightpaths are costly, active and transfer learning techniques can be used to reduce both the number of probing lightpaths and the dataset size required for training QoT models of sufficient accuracies [22].

It is worth mentioning, that the loss function of Eq. (1) does not incur any additional computational complexity during DNN model training; rather, the computational complexity of the model is only affected by the DNN architecture (i.e., number of hidden layers, number of units, activation functions considered) and the optimization algorithm applied for training the DNN model. Hence, the computational complexity of the q-quantile DNN model is similar to the complexity of a DNN model trained according to the commonly considered least squares loss function (i.e., the MSE loss function), as long as the underlying DNN model architecture is the same. A theoretical analysis regarding the computational efficiency of training neural networks can be found in [39].

3. Dataset generation

Dataset D was generated on the US backbone network (USNET) topology (Fig. 1) assuming an EON with single-mode single-core fibers operating with BPSK, QPSK, 8-QAM, and 16-QAM modulation formats. Each fiber link has a capacity of 320 frequency slots (FSs), spaced at 12.5 GHz, and with a baud rate of 10.5 Gbaud. Five thousand connection requests were generated in a dynamic network following a Poisson process with exponentially distributed holding times for a network load equal to 300 Erlangs. Each connection request, defined by a source node, a destination node, and a bit-rate demand, was generated as follows: a source–destination pair was randomly sampled from the set of network nodes and a bit-rate demand was randomly sampled from the interval $[50, 200]$ Gbps following a uniform distribution. For the connection provisioning phase, a conventional RSA algorithm was implemented during which dataset D was created by extracting the input and output features from each computed lightpath.

Specifically, for each arriving connection request, the routing sub-problem was solved according to the k-shortest path algorithm [40], with $k = 3$, while the spectrum allocation (SA) sub-problem was solved according to the first-fit scheme, ensuring that all three SA constraints were met; that is, the spectrum continuity, contiguity, and no frequency overlap constraints. Note that for each connection request, the spectrum demand, measured in frequency slots, was computed according

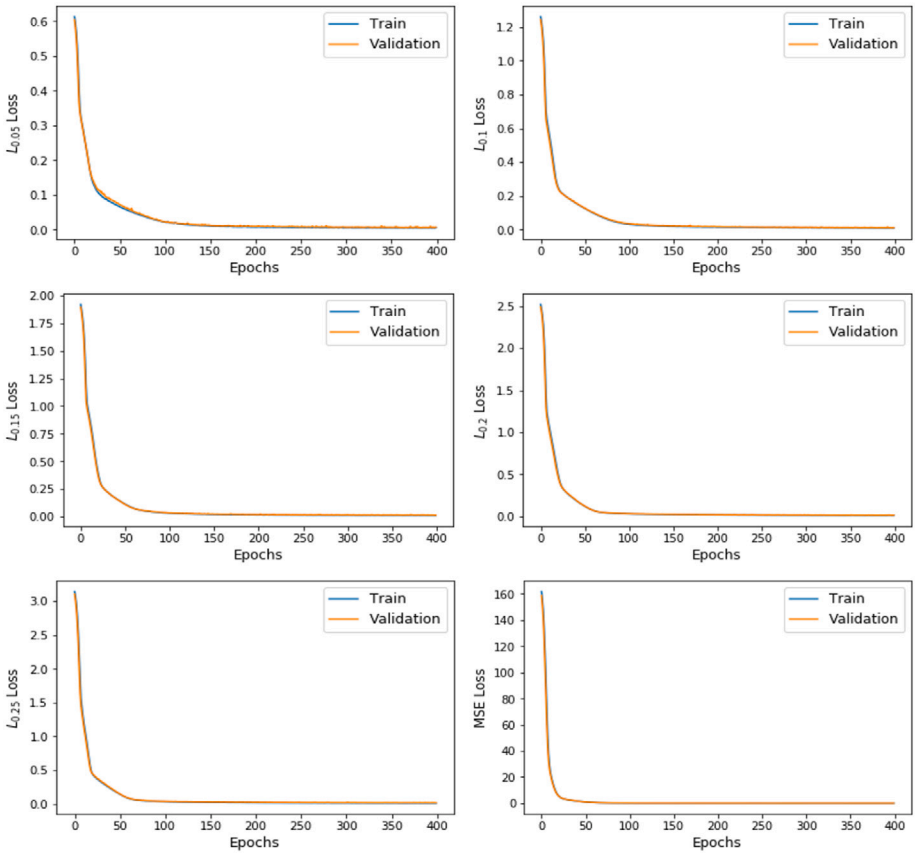

Fig. 2. Loss versus the number of training epochs for various loss functions: $L_{0.05}$, $L_{0.1}$, $L_{0.15}$, $L_{0.2}$, $L_{0.25}$, MSE loss functions.

to a conventional distance-adaptive modulation scheme. Specifically, a modulation format was selected according to the physical distance of the selected route, and the number of requested FSs was computed according to the bit-rate demand [41]. A connection was considered as blocked (not provisioned) if a feasible spectrum allocation could not be found for any of the k-shortest paths computed for that connection. For each provisioned lightpath i, the features described in Section 2 were extracted to form input patterns $\mathbf{x}_i$, while for the output QoT values, y_i, the Q-tool analytically described in [24] was used. Hence, our dataset D was synthetically generated, following a common approach in the literature when real datasets are not available [23]. As 349 connection requests were blocked, resulting to 7% blocking probability, $n = 4651$ patterns were included in dataset D.

4. Model training and evaluation

Dataset D was used for training several QoT models, with each QoT model minimizing a different loss function. Specifically, for quantile regression, the loss function in Eq. (1) was applied for $q = 0.05$, $q = 0.1$, $q = 0.15$, $q = 0.2$, and $q = 0.25$ to create the lower QoT estimates, denoted as $\hat{y}_*^q$, for any unseen lightpath pattern $\mathbf{x}_*$ (i.e., a lightpath pattern that was not used during training). The different q-quantile models serve to demonstrate the importance of fine-tuning the q values towards network efficient and confident decision-making, as discussed in the sections that follow. Additionally, a QoT model was trained to minimize the mean squared error (MSE) loss function to generate the QoT estimate (i.e., prediction), denoted as $\hat{y}_*^p$ in the conventional way; that is, these predictions correspond to the mean response of the model, commonly considered in the QoT estimation literature. In this work, these predictions are utilized to estimate traditional empirical margins and perform decisions in the conventional way, hence allowing a comparison with the proposed approach (i.e., used as benchmarks).

All quantile and least squares QoT models were parameterized by a DNN model with 3 hidden layers of 64, 32, and 32 hidden units, respectively. The rectified linear unit (ReLU) activation function [42] was used for all hidden units. Training was performed according to the efficient, stochastic gradient decent (SGD), Adam optimization algorithm [43] for 400 epochs, with a batch size equal to 50 and a learning rate equal to 10^{-4}. Before training, dataset D was scaled in order for the input features to follow the standard normal distribution. Seventy percent (70%) of the patterns in D were used for training, from which twenty percent (20%) were used for validation purposes. Hence, thirty percent (30%) of the patterns in D were used for testing (i.e., 1395 test patterns in total).

Fig. 2 illustrates $L_{0.0.5}$, $L_{0.01}$, $L_{0.15}$, $L_{0.2}$, $L_{0.25}$, and MSE loss performance over the training epochs. Note that for training each model according to all 400 epochs, approximately up to 7 seconds of training time was required. Clearly, the loss for both training and validation datasets reduces and eventually approaches zero as the number of epochs increases for all loss functions considered. Performance accuracy in the test dataset closely follows the accuracy obtained in the validation dataset for all models. Importantly, to validate the accuracy of all quantile models with respect to the percentage of test patterns that fall below the considered q-quantile, all $\hat{y}_*^q$ estimates were compared with their ground truths y_*.

Table 1 summarizes the validation accuracy of all quantile QoT models. According to Table 1, all q-quantiles are sufficiently approximated as the validation accuracy is close to the pre-selected q value. Indicatively, the 0.05-quantile obtains lower estimation bounds for 93% of the test patterns, with only 7% of the estimates falling below their true QoT value. Note that, even though a perfect q-quantile model can theoretically achieve $q \times 100\%$ of patterns falling below their true QoT value, slight variations between the q value and the validation accuracy is expected, since the q-quantiles are just DNN approximations.

Table 1
Quantile QoT model validation.

Value of q	0.05	0.1	0.15	0.2	0.25
Validation accuracy	0.07	0.12	0.14	0.21	0.24

Table 2
Estimated margins (in dB).

M_w	M_e	$\bar{M}_{0.05} \pm std$	$\bar{M}_{0.1} \pm std$	$\bar{M}_{0.15} \pm std$	$\bar{M}_{0.2} \pm std$	$\bar{M}_{0.25} \pm std$
0.52	0.35	0.15 ± 0.1	0.13 ± 0.1	0.11 ± 0.1	0.09 ± 0.09	0.08 ± 0.08

Nevertheless, to sufficiently approximate a q-quantile model, several models with diverse hyperparameters can be trained and the one that results in the best validation accuracy within the test dataset can be opted for. It should be noted that this procedure does not hinder the practical applicability of the overall framework, since candidate DNN models can be trained off-line and in parallel within seconds. In this work, all models required less than 7 seconds to converge.

5. Margin estimation and comparisons

A design margin is typically considered to compensate for the inaccuracies of ML-aided QoT models, usually optimized according to least squares techniques (e.g., optimized to minimize the MSE loss function). Specifically, a design margin is appropriately estimated and used to address the inaccuracies resulting from the predictions overestimating the true QoT of lightpaths. Since the predictions that will end up overestimating the true QoT of unseen lightpaths cannot be known a-priori, this margin is used to penalize the predictions of all unseen lightpaths. For margin estimation purposes, several approaches have been proposed, such as the worst error margin [6] and the empirical rule error margin [5]. In this work, both these approaches are used as benchmarks. In this section, both margins are initially defined, and then their corresponding values are estimated according to the QoT model trained to minimize the MSE loss function in Section 4. The equivalent average quantile margin metric is also defined, that is subsequently used for comparison purposes; that is, to examine whether the q-quantiles result in margin reduction, as opposed to the existing margin estimation approaches.

5.1. Worst error margin

The worst error margin (M_w) is defined as the maximum error resulting between the true and predicted QoT values amongst all lightpaths with their true QoT overestimated by the predictions:

$$M_w = \max_{y_i < \hat{y}_i^p} \{|y_i - \hat{y}_i^p|\}_{i=1}^{n'}, \tag{2}$$

where n' is the number of patterns in the test dataset, and $\hat{y}_i^p$ is the QoT prediction (i.e., mean response of the model).

5.2. Empirical rule error margin

The empirical rule error margin (M_e) is based on the empirical rule error stating that 99.7% of the data (i.e., errors) will be within three standard deviations of the mean error (i.e., $\mu \pm 3\sigma$), provided that the errors follow a Gaussian distribution. Hence, to account for the worst error, M_e is given by:

$$M_e = \mu + 3\sigma, \tag{3}$$

where $\mu = \frac{1}{n'} \sum_{i=1}^{n'} |y_i - \hat{y}_i^p|$ is the mean absolute error resulting between the true and predicted QoT values in the test dataset, and $\sigma = \sqrt{\frac{1}{n'} \sum_{i=1}^{n'} (y_i - \hat{y}_i^p)^2}$ is the standard deviation of the error.

Table 3
Margin reduction (in dB).

	$q = 0.05$	$q = 0.1$	$q = 0.15$	$q = 0.2$	$q = 0.25$
$\bar{R}_{q\|w} \pm std$	71.6 ± 19.6	76.2 ± 20.4	78.3 ± 19.8	81.8 ± 18.3	83.6 ± 16.4
$\bar{R}_{q\|e} \pm std$	57.1 ± 29.5	64.17 ± 30.8	67.2 ± 29.8	72.3 ± 27.6	75.3 ± 24.8

5.3. Equivalent average quantile margin

In the proposed framework, the q-quantile model automatically predicts a lower QoT estimate for each individual lightpath. Hence, in practice, margin estimation is not required, as the outputs of the lower quantile model can be directly used for decision-making. However, for comparison purposes, to evaluate whether the proposed approach is successful in reducing the margins considered in the benchmark approaches, the quantile equivalent margin ($\bar{M}_q$) is defined as the mean difference between the predicted and the q-quantile QoT estimates of all lightpaths in the test dataset. Specifically,

$$\bar{M}_q = \frac{1}{n'} \sum_{i=1}^{n'} (\hat{y}_i^p - \hat{y}_i^q). \tag{4}$$

Hence, just like M_w and M_e, an M_q value, computed for a single lightpath pattern, indicates the equivalent penalty applied to a least squares prediction (i.e., indicates by how much a prediction is reduced to reach the lower q-quantile value considered for decision-making for a single lightpath). Since the $M_{q|i}$ value is different for each lightpath i (i.e., $M_{q|i} = \hat{y}_i^p - \hat{y}_i^q$), the average quantile margin $\bar{M}_q$ is computed over all test lightpaths.

5.4. Margin comparisons

All types of margins are estimated according to the trained models and corresponding test dataset of Section 4 and summarized in Table 2. Clearly, $\bar{M}_q$ outperforms both benchmarks for all q values considered. Interestingly, as q increases, $\bar{M}_q$ further reduces. This trend is reasonable, since higher q values increase the percentage of unseen patterns that may fall below the q-quantile, relaxing the lower QoT estimation bounds and consequently leading to further margin reduction.

To quantify the achievable margin reduction, $\bar{R}_{q|*}$ is defined as the average quantile margin reduction over margin M_*, where $*$ denotes either the worst error margin (i.e., w index) or the empirical rule error margin (i.e., e index). Hence, the average quantile margin reduction is given by:

$$\bar{R}_{q|*} = \frac{100}{n'} \sum_{i=1}^{n'} \frac{(M_* - M_{q|i})}{M_*}. \tag{5}$$

Table 3 summarizes the margin reduction achieved for all q values, when compared to the benchmark margins. According to these results, margin reduction varies between 71% to 83% when the quantile margin is compared to M_w and between 57% to 75% when compared to M_e. The highest reduction is observed for the highest q value examined. Note that, even though in this work M_e outperforms M_w, in the sense that M_e is lower than M_w, this may not always be the case, since their values merely depend on the dataset available and on what can be observed from the test patterns (i.e., it is possible that in the available test dataset the highest possible deviation between the true and predicted patterns is not observed). On the other hand, the behavior of the q-quantiles tends to be stable, since the quantile margins are not estimated according to the test patterns, but rather they are learned during training and inferred for each unseen lightpath.

Insights on why q-quantiles achieve such a reduction as q increases are given by observing Fig. 3. Fig. 3 illustrates the quantile estimates $\hat{y}_*^q$ for $q = 0.05$ and $q = 0.25$, the mean response predictions $\hat{y}_*^p$, and the ground truths y_*, for a small number of lightpath samples randomly selected from the test dataset. Clearly, the ground truths of some

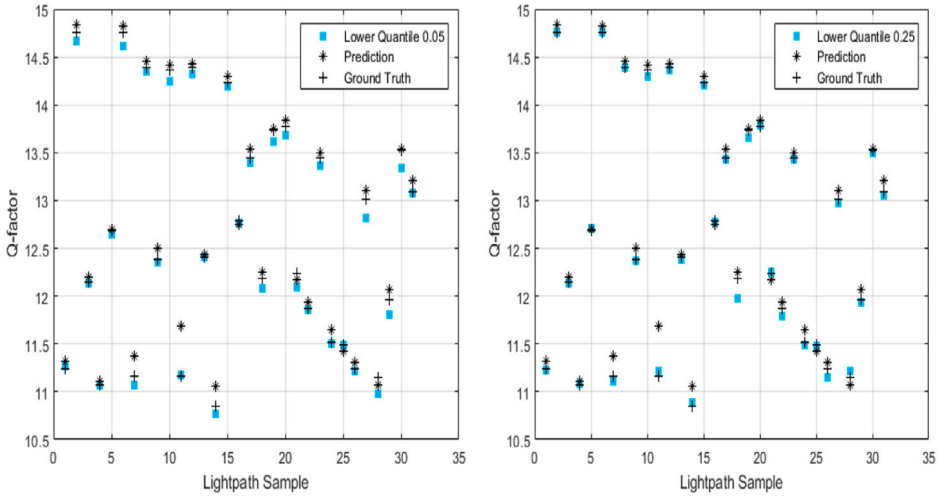

Fig. 3. Ground truths, predictions, and quantile estimates for $q = 0.05$ and $q = 0.25$, over a number of lightpath samples.

lightpaths are tightly bounded by the lower quantiles (i.e., the quantile estimates and the ground truths are close to each other), while for other lightpaths the deviation between lower quantiles and ground truths is larger. In general, the deviation between lower quantiles and ground truths varies, depending on the uncertainty level of the q-quantile model, over each unseen lightpath. The lower quantiles capture exactly this fact by discriminating between the different inputs, hence returning a QoT estimate that is not lower than necessary, given the desired certainty level (i.e., the q value). Importantly, the deviation between lower quantiles and ground truths seems to be mitigated when $q = 0.25$, as compared to the case where $q = 0.05$, demonstrating the importance of fine-tuning the q value prior to decision-making.

Unlike quantile margins, the M_w and M_e margins cannot distinguish between input patterns, hence they unnecessarily penalize all lightpaths. As an example, the M_w margin will unnecessarily reduce the predictions of several lightpaths (e.g., samples $1, 3, 5, 7, 12$ in Fig. 3 for which the MSE model is almost certain) by a large constant margin to mitigate the error of sample 11 (i.e., sample 11 results in the worst error margin M_w). The same holds for the M_e margin, which is lower than M_w, but will still unnecessarily reduce the predictions of these lightpaths. On the other hand, lower quantiles for these lightpaths suggest that reductions of such magnitude are not required, with high certainty.

Nevertheless, lower quantile estimates still have a probability of failing, depending on the q-quantile selected. This is clearly shown in sample 28 of Fig. 3, where its ground truth falls below the 0.25 quantile estimate. Note that this situation does not appear when the 0.05-quantile is considered, at least for the samples illustrated in Fig. 3. Clearly, even though for both quantiles there exists a probability for this situation to occur, this probability reduces to zero as $q \to 0$. However, in decision-making, pushing this probability to zero is not of the utmost importance; instead, it is more important to sufficiently tune this probability to the point where decision accuracy, in both classes of interest, is the best possible; that is, the probability to correctly recognize lightpaths with insufficient QoT is close to 1, without significantly reducing the probability to correctly recognize lightpaths with sufficient QoT.

6. Decision accuracy and comparisons

To investigate the impact of each margin estimation approach on decision accuracy, the two classes of interest according to a predefined QoT threshold are initially defined, and then the outputs of each QoT model are compared against this threshold. Specifically, each unseen lightpath is classified as infeasible if its QoT estimate is below the QoT threshold (Class 1), otherwise it is classified as feasible (Class 2). Note that the estimates of the least squares QoT models are first reduced either by the M_e or M_w margin before classification takes place, depending on the margin scheme applied. The estimates of the q-quantile QoT models are directly used for classification purposes.

Table 4
Classification accuracy for several QoT thresholds.

	Threshold 12 dB						
	M_w	M_e	$M_{0.05}$	$M_{0.1}$	$\mathbf{M_{0.15}}$	$M_{0.2}$	$M_{0.25}$
Accuracy	0.86	0.92	0.98	0.98	**0.98**	0.98	0.99
Class 1 accuracy	1	1	0.997	0.995	**0.997**	0.984	0.995
Class 2 accuracy	0.8	0.86	0.98	0.98	**0.97**	0.98	0.98
	Threshold 12.5 dB						
	M_w	M_e	$M_{0.05}$	$M_{0.1}$	$\mathbf{M_{0.15}}$	$M_{0.2}$	$M_{0.25}$
Accuracy	0.84	0.89	0.96	0.97	**0.97**	0.98	0.98
Class 1 accuracy	1	1	0.998	0.995	**1**	0.995	0.997
Class 2 accuracy	0.73	0.82	0.94	0.94	**0.95**	0.96	0.98
	Threshold 13 dB						
	M_w	M_e	$M_{0.05}$	$M_{0.1}$	$M_{0.15}$	$\mathbf{M_{0.2}}$	$M_{0.25}$
Accuracy	0.88	0.91	0.96	0.96	0.97	**0.98**	0.98
Class 1 accuracy	1	1	0.997	1	0.995	**0.997**	0.996
Class 2 accuracy	0.7	0.78	0.89	0.91	0.94	**0.94**	0.96
	Threshold 13.5 dB						
	M_w	M_e	$M_{0.05}$	$M_{0.1}$	$M_{0.15}$	$\mathbf{M_{0.2}}$	$M_{0.25}$
Accuracy	0.9	0.94	0.97	0.97	0.98	**0.98**	0.98
Class 1 accuracy	1	1	0.999	0.998	0.999	**1**	0.997
Class 2 accuracy	0.65	0.78	0.9	0.91	0.94	**0.95**	0.94
	Threshold 14 dB						
	M_w	M_e	$M_{0.05}$	$M_{0.1}$	$\mathbf{M_{0.15}}$	$M_{0.2}$	$M_{0.25}$
Accuracy	0.9	0.95	0.98	0.98	**0.99**	0.99	0.99
Class 1 accuracy	1	1	1	1	**1**	0.999	0.998
Class 2 accuracy	0.44	0.7	0.9	0.91	**0.94**	0.96	0.97

As shown in Table 4, several QoT thresholds are considered, with the classification accuracy evaluated on the lightpaths of the test dataset. Hence, decisions are taken according to the inferred QoT values reduced by the margin considered (if any, depending on the margin scheme applied), and accuracy is tested against the ground truths in the test dataset. Table 4 denotes each approach according to the corresponding margin scheme applied; that is, M_w, M_e, and M_q for the various q values considered.

According to the results of Table 4, both the M_w and M_e schemes achieve 100% accuracy for Class 1, hence ensuring that a lightpath with insufficient QoT will never be established in the network. However, the

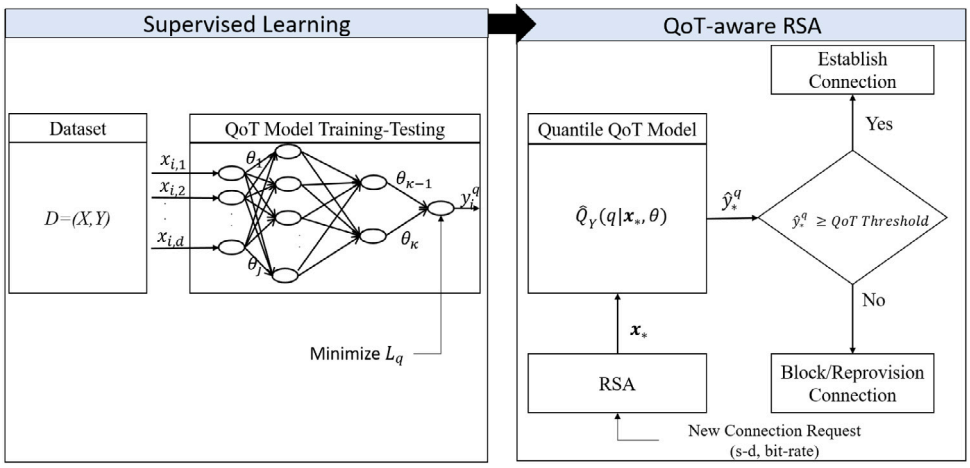

Fig. 4. General framework of the quantile QoT model training and QoT-aware RSA over unseen lightpaths.

drawback of these schemes is that they demonstrate a low accuracy for Class 2 (varying between 44% to 86%), which is a direct consequence of the fact that these margins tend to greatly underestimate the true QoT of all lightpaths. In contrast, the quantile approach, for all M_q values considered, similarly succeeds for Class 1, while also achieving a high classification accuracy for Class 2 (varying between 94% to 98%). It should be noted that a slight loss in classification accuracy of Class 1 is observed (up to 1.6% for the q-quantile value performing the worst, i.e., for $q = 0.2$ and QoT threshold equal to 12 dB), due to the significant margin reduction that quantiles achieve when compared to the M_w and M_e values.

Nevertheless, the results clearly indicate the importance of appropriately fine-tuning the q values to achieve a negligible loss for Class 1, while also preserving a high accuracy for Class 2. Specifically, according to Table 4, this loss varies between 0% to 0.3% by appropriately selecting the q-quantiles performing the best for each QoT threshold considered (shown in bold in Table 4), with the improvement in Class 2 accuracy ranging between 11% to 25% when 0.3% loss is observed. Overall, however, the improvement for Class 2 can be up to 53% when q-quantile models are opted for, depending on the QoT threshold considered. In general, to identify a good q-quantile, a straightforward approach is to examine various incremental q values up to the point where the accuracy for Class 1 is considered sufficient. Identifying an optimal q value in an automated manner is planned for future work.

7. Network performance evaluation

To further demonstrate the importance of closely capturing uncertainty prior to QoT related decisions (i.e., fine-tuning margins over lightpath uncertainty), the impact of the various approaches on network performance (in terms of blocking probability) was also examined according to a dynamic QoT-aware RSA framework. Specifically, in the dynamic QoT-aware RSA framework, decisions regarding the QoT feasibility of arriving connection requests were taken according to one of the various DNN-QoT models and margin estimation approaches examined in this work; that is,

- QoT related decisions were taken according to a least squares DNN-QoT model with the QoT estimates degraded according to the state-of-the-art M_w margin before decision-making,
- QoT related decisions were taken according to a least squares DNN-QoT model with the QoT estimates degraded according to the state-of-the-art M_e margin before decision-making,
- QoT related decisions were taken according to the estimates of the proposed q-quantile DNN-QoT model. Note that the estimates of the q-quantile DNN-QoT model are directly considered for decision-making, without any additional margin reduction, as they already account for margins during inference.

Additionally, we compare all approaches against the decisions taken by the Q-tool used for generating dataset D. Hence, the Q-tool results serve to examine how close each margin estimation approach is to the accurate decisions in terms of blocking probability.

The general framework of the quantile-based dynamic QoT-aware RSA scheme is described in Fig. 4, illustrating that quantile QoT model training/testing is performed off-line, by means of supervised learning, according to a dataset D already collected from a dynamic network. Once the quantile QoT model is available, the QoT-aware RSA is executed on-line for each new connection request, described by a source–destination (s-d) pair and a bit-rate demand, to extract input pattern $\mathbf{x}_*$ of the unseen (new) lightpath. Specifically, for each $\mathbf{x}_*$, the quantile QoT model returns the lower quantile estimate $\hat{y}_*^q$ which is compared against a predefined QoT threshold. The unseen lightpath is established if $\hat{y}_*^q$ is equal to or above the acceptable QoT threshold; otherwise, it is blocked or reprovisioned over a different lightpath. In our work, for simplicity, we opted to directly block the connection if the QoT of a lightpath is not sufficient. Note that the dynamic QoT-aware RSA framework of Fig. 4 is similar for every margin estimation scheme examined in this work, with the difference that a different loss function is minimized during QoT model training when the baseline schemes are considered (i.e., the MSE loss function), and that after model inference a precomputed margin (i.e., M_w or M_e) is used to degrade $\hat{y}_*$ QoT estimates before comparing them against the QoT threshold.

Hence, for the simulations, for each margin estimation approach, 6000 connection requests were generated considering the EON parameters of Section 3 for a network load of 200 Erlangs and 160 FSs on each network link. For the connection provisioning phase a conventional QoT-aware RSA algorithm was implemented, in which a lightpath was considered feasible if both the SA and QoT constraints were met. Lightpath feasibility was examined by iterating over the k-shortest paths (with $k = 3$), and a connection was blocked if a feasible SA with a QoT that is above the QoT threshold could not be found.

For the first 5000 connection requests, QoT feasibility was examined according to the Q-tool described in [24], while a dataset D was at the same time created for training and testing the QoT model under consideration. Hence, after the first 5000 connection requests, a DNN-QoT model was trained to replace the Q-tool for the next 1000 connections. Depending on the margin estimation approach followed, M_w, M_e, or a margin equal to zero was computed and used to degrade the DNN-QoT model estimates before examining the QoT feasibility of each computed lightpath.

For the DNN-QoT models trained to minimize the MSE loss function, M_w or M_e was computed. For the DNN-QoT model trained to minimize a q-quantile loss function, the margin was set to zero. Note that for the latter approach, the DNN-QoT model was validated not only according

Table 5
Network performance evaluation with QoT threshold equal to 12 dB.

	Q-tool	DNN-QoT with M_w	DNN-QoT with M_e	DNN-QoT with $M_{0.15}$
Blocking probability	0.28	0.37	0.35	0.3
Blocking probability due to QoT	0.25	0.34	0.31	0.26
Blocking probability due to capacity	0.03	0.03	0.04	0.04

to its performance over the test patterns but also according to the accuracy of the percentage of test patterns that fell below the lower q-quantile estimate (i.e., according to the accuracy with respect to the q value). All DNN-QoT models were trained according to the parameters described in Section 4, with the number of epochs increased to 500 to account for the higher diversity of patterns in the newly created datasets, that is a direct consequence of the added QoT constraint which is considered in the QoT-aware RSA. Training time for each model was on the order of seconds (i.e., up to 7 seconds in the worst case for the MSE QoT model), while inference time was on the order of milliseconds, enabling the on-line provisioning and QoT-related decisions to be performed within an acceptable time (i.e., appropriate for real implementations).

Table 5 summarizes the average results obtained from a set of four simulation runs for each different approach examined. The QoT threshold was set to 12 dB which corresponds to the case where the baseline schemes performed the best with respect to the achievable classification accuracy in both classes of interest (Table 4). Each approach in Table 5 is denoted according to the margin estimation approach followed (i.e., DNN-QoT with M_w, DNN-QoT with M_e, and DNN-QoT with M_q). Note that q is set to 0.15, which corresponds to the q value performing the best for the 12 dB QoT threshold according to the results in Table 4. Furthermore, Table 5 reports the average results obtained from the Q-tool from which dataset D was created for training the models; that is, the Q-tool results correspond to the blocking probability obtained for the first 5000 connection requests and the DNN-QoT results for all margin approaches correspond to the blocking probability obtained for the last 1000 connection requests.

Overall, Table 5 clearly indicates that the q-quantile approach (i.e., DNN-QoT with $M_{0.15}$) performs closer to the results obtained with the Q-tool, when compared to the baseline techniques. This is to be expected, since the q-quantile estimates tend to underestimate less the true QoT of lightpaths when compared to the baseline approaches, hence leading to more accurate QoT-related decisions. Importantly, the q-quantile scheme achieves a blocking probability that is improved approximately by 14% when compared to the DNN-QoT with M_e scheme and by 19% when compared to the DNN-QoT with M_w scheme. Clearly, fine-tuning margins over lightpath uncertainty has a significant impact on the reduction of blocking probability, and deep quantile regression constitutes a promising approach towards this objective.

To gain further insights on the reasons why the baseline schemes experience a higher blocking probability compared to the q-quantile approach, Table 5 reports also the percentages of blocking caused due the QoT constraint and due to the capacity constraints (i.e., the SA constraints). Evidently, this higher blocking probability is mainly caused due to the QoT constraint as a consequence of the conservative margins that baseline approaches consider (with higher blocking probability when M_w is considered, as the M_e margin tends to be lower than the M_w margin). Clearly, the lowest average margin is achieved by the q-quantile scheme with $q = 0.15$, that only slightly increases the blocking probability due to QoT (i.e., by 3%) when compared to the accurate QoT-related decisions (i.e., the decisions derived by the Q-tool). This slight increase is to be expected, since the q-quantiles still experience inaccuracies in the class of lightpaths with sufficient QoT (i.e., 3% misclassification for Class 2 according to Table 4 for a QoT threshold of 12 dB and $M_{0.15}$).

It should be noted that for these simulation results the QoT threshold performing the best for the baseline schemes was considered; thus, it is expected that the improvement for the rest of the QoT thresholds will be even higher. Further, it is worth mentioning that while the datasets of this work were synthetically generated, this does not affect the scope of this work, which is to demonstrate the potential of quantile regression to better capture uncertainty over unseen lightpaths. As the uncertainty is expected to be higher in real datasets (e.g., due to noisy OPM inputs), such frameworks are expected to be even more important in real network implementations. Finally, more accurately capturing uncertainty for other network operation decisions (e.g., ML-aided traffic-driven SA decisions to reduce over-provisioning by fine-tuning margin estimations [44]) is expected to further improve network performance; that is, the cumulative improvement on network performance is expected to be significant when uncertainty is more accurately considered for all data-driven network operations.

8. Conclusion

This work demonstrates the capabilities of deep quantile regression for sufficiently addressing uncertainty of QoT estimates over unseen lightpaths. The main advantage of this framework is that uncertainty is addressed through the discriminative inference of margins over the unseen lightpaths, alleviating the need to resort to empirically estimated myopic margins. It is shown that q-quantile DNN-QoT models achieve, on average, significant margin reduction (up to 83%) when compared to baseline margin estimation approaches. Importantly, it is also shown that decisions based on q-quantile DNN-QoT models achieve classification accuracies that are high for both classes of interest. Specifically, q-quantiles succeed with almost 100% accuracy to identify lightpaths with insufficient QoT, while preserving a high accuracy (i.e., above 91%) for the class of feasible lightpaths. In contrast, baseline approaches succeed only for the class of infeasible lightpaths, significantly increasing the error for the class of feasible lightpaths (up to 65% error). While a small error is observed in the critical class of infeasible lightpaths when a q-quantile is opted for (up to 1.6% error), this can be further reduced (i.e., to 0.3% error) by fine-tuning the q value considered for model training. In any case, errors of such magnitude are negligible, considering the fact that the improvement for the class of feasible lightpaths is up to 53% when compared to the baseline techniques. As the promise of margin reduction and accurate decision-making is ultimately to improve network performance, it is also shown that the q-quantiles are capable of achieving significant improvements (up to 19%) in terms of blocking probability, when compared to the baseline margin estimation schemes.

Examining and comparing various ML methods capable of addressing model uncertainty (e.g., MC dropout) is planned as future work.

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 European Union's Horizon 2020 research and innovation programme under grant agreement No 739551 (KIOS CoE - TEAMING) and from the Republic of Cyprus through the Deputy Ministry of Research, Innovation and Digital Policy.

References

[1] T. Panayiotou, S.P. Chatzis, G. Ellinas, Performance analysis of a data-driven quality-of-transmission decision approach on a dynamic multicast- capable metro optical network, IEEE/OSA J. Opt. Commun. Networking 9 (1) (2017) 98–108.

[2] R. Proietti, X. Chen, K. Zhang, G. Liu, M. Shamsabardeh, A. Castro, L. Velasco, Z. Zhu, S.J. Ben Yoo, Experimental demonstration of machine-learning-aided QoT estimation in multi-domain elastic optical networks with alien wavelengths, IEEE/OSA J. Opt. Commun. Networking 11 (1) (2019) A1–A10.

[3] R.M. Morais, J. Pedro, Machine learning models for estimating quality of transmission in DWDM networks, IEEE/OSA J. Opt. Commun. Networking 10 (10) (2018) D84–D99.

[4] G. Bergk, B. Shariati, P. Safari, J.K. Fischer, ML-assisted QoT estimation: A dataset collection and data visualization for dataset quality evaluation, IEEE/OSA J. Opt. Commun. Networking 14 (3) (2022) 43–55.

[5] E. Seve, J. Pesic, C. Delezoide, S. Bigo, Y. Pointurier, Learning process for reducing uncertainties on network parameters and design margins, IEEE/OSA J. Opt. Commun. Networking 10 (2) (2018) A298–A306.

[6] I. Sartzetakis, K.K. Christodoulopoulos, E.M. Varvarigos, Accurate quality of transmission estimation with machine learning, IEEE/OSA J. Opt. Commun. Networking 11 (3) (2019) 140–150.

[7] Y. Gal, Z. Ghahramani, Dropout as a Bayesian approximation: Representing model uncertainty in deep learning, in: Proc. 33rd International Conference on International Conference on Machine Learning (ICML), 2016, pp. 1050–1059.

[8] F. Rodrigues, F.C. Pereira, Beyond expectation: Deep joint mean and quantile regression for spatiotemporal problems, IEEE Trans. Neural Netw. Learn. Syst. 31 (12) (2020) 5377–5389.

[9] R. Koenker, Fundamentals of quantile regression, in: Econometric Society Monographs, Cambridge University Press, 2005, pp. 26–67.

[10] J. Mata, I. de Miguel, R.J. Durán, N. Merayo, S.K. Singh, A. Jukan, M. Chamania, Artificial intelligence (AI) methods in optical networks: A comprehensive survey, Opt. Switch. Netw. 28 (2018) 43–57.

[11] F. Musumeci, C. Rottondi, A. Nag, I. Macaluso, D. Zibar, M. Ruffini, M. Tornatore, An overview on application of machine learning techniques in optical networks, IEEE Commun. Surv. Tutor. 21 (2) (2019) 1383–1408.

[12] R. Boutaba, M.A. Salahuddin, N. Limam, S. Ayoubi, N. Shahriar, F.E. Solano, O.M. Caicedo, A comprehensive survey on machine learning for networking: Evolution, applications and research opportunities, J. Internet Serv. Appl. 9 (16) (2018) 1–99.

[13] X.-L. Huang, X. Ma, F. Hu, Editorial: Machine learning and intelligent communications, Mobile Netw. Appl. 23 (1) (2018) 68–70.

[14] T. Panayiotou, K. Manousakis, S.P. Chatzis, G. Ellinas, A data-driven bandwidth allocation framework with QoS considerations for EONs, IEEE/OSA J. Lightwave Technol. 37 (9) (2019) 1853–1864.

[15] X. Chen, R. Proietti, H. Lu, A. Castro, S.J.B. Yoo, Knowledge-based autonomous service provisioning in multi-domain elastic optical networks, IEEE Commun. Mag. 56 (8) (2018) 152–158.

[16] R. Alvizu, S. Troia, G. Maier, A. Pattavina, Matheuristic with machine-learning-based prediction for software-defined mobile metro-core networks, IEEE/OSA J. Opt. Commun. Networking 9 (9) (2017) D19–D30.

[17] T. Panayiotou, S.P. Chatzis, G. Ellinas, Leveraging statistical machine learning to address failure localization in optical networks, IEEE/OSA J. Opt. Commun. Networking 10 (3) (2018) 162–173.

[18] B. Shariati, M. Ruiz, J. Comellas, L. Velasco, Learning from the optical spectrum: Failure detection and identification, IEEE/OSA J. Lightwave Technol. 37 (2) (2019) 433–440.

[19] X. Chen, C.-Y. Liu, R. Proietti, J. Yin, Z. Li, S.J.B. Yoo, On cooperative fault management in multi-domain optical networks using hybrid learning, IEEE J. Sel. Top. Quantum Electron. (2022) 1.

[20] C. Natalino, M. Schiano, A. Di Giglio, L. Wosinska, M. Furdek, Field demonstration of machine-learning-aided detection and identification of jamming attacks in optical networks, in: Proc. European Conference on Optical Communication (ECOC), 2018, pp. 1–3.

[21] M. Bensalem, S.K. Singh, A. Jukan, On detecting and preventing jamming attacks with machine learning in optical networks, in: Proc. IEEE Global Communications Conference (GLOBECOM), 2019, pp. 1–6.

[22] D. Azzimonti, C. Rottondi, A. Giusti, M. Tornatore, A. Bianco, Comparison of domain adaptation and active learning techniques for quality of transmission estimation with small-sized training datasets, IEEE/OSA J. Opt. Commun. Networking 13 (1) (2021) A56–A66.

[23] Y. Pointurier, Machine learning techniques for quality of transmission estimation in optical networks, IEEE/OSA J. Opt. Commun. Networking 13 (4) (2021) B60–B71.

[24] B. Shariati, A. Mastropaolo, N.-P. Diamantopoulos, J.M. Rivas-Moscoso, D. Klonidis, I. Tomkos, Physical-layer-aware performance evaluation of SDM networks based on SMF bundles, MCFs, and FMFs, IEEE/OSA J. Opt. Commun. Networking 10 (9) (2018) 712–722.

[25] A. Carena, G. Bosco, V. Curri, Y. Jiang, P. Poggiolini, F. Forghieri, EGN model of non-linear fiber propagation, OSA Opt. Express 22 (13) (2014) 16335–16362.

[26] G. Ellinas, N. Antoniades, T. Panayiotou, A. Hadjiantonis, A.M. Levine, Multicast routing algorithms based on Q-factor physical-layer constraints in metro networks, IEEE Photonics Technol. Lett. 21 (6) (2009) 365–367.

[27] S. Yan, F.N. Khan, A. Mavromatis, D. Gkounis, Q. Fan, F. Ntavou, K. Nikolovgenis, F. Meng, E.H. Salas, C. Guo, C. Lu, A.P.T. Lau, R. Nejabati, D. Simeonidou, Field trial of machine-learning-assisted and SDN-based optical network planning with network-scale monitoring database, in: Proc. European Conference on Optical Communication (ECOC), 2017, pp. 1–3.

[28] T. Panayiotou, G. Savva, B. Shariati, I. Tomkos, G. Ellinas, Machine learning for QoT estimation of unseen optical network states, in: Proc. IEEE/OSA Optical Fiber Communication Conference (OFC), paper Tu2E.2, 2019, pp. 1–3.

[29] G. Savva, T. Panayiotou, I. Tomkos, G. Ellinas, Deep graph learning for QoT estimation of unseen optical sub-network states: Capturing the crosstalk impact on the in-service lightpaths, IEEE/OSA J. Lightwave Technol. 40 (4) (2022) 921–934.

[30] T. Panayiotou, G. Savva, I. Tomkos, G. Ellinas, Decentralizing machine-learning-based QoT estimation for sliceable optical networks, IEEE/OSA J. Opt. Commun. Networking 12 (7) (2020) 146–162.

[31] L. Zhang, X. Li, Y. Tang, J. Xin, S. Huang, A survey on QoT prediction using machine learning in optical networks, Opt. Fiber Technol., Mater. Devices Syst. 68 (102804) (2022).

[32] Z. Gao, S. Yan, J. Zhang, M. Mascarenhas, R. Nejabati, Y. Ji, D. Simeonidou, ANN-based multi-channel QoT-prediction over a 563.4-km field-trial testbed, IEEE/OSA J. Lightwave Technol. 38 (9) (2020) 2646–2655.

[33] P. Paudyal, S. Shen, S. Yan, D. Simeonidou, Toward deployments of ML applications in optical networks, IEEE Photonics Technol. Lett. 33 (11) (2021) 537–540.

[34] X. Chen, B. Li, R. Proietti, C.-Y. Liu, Z. Zhu, S.J.B. Yoo, Demonstration of distributed collaborative learning with end-to-end QoT estimation in multi-domain elastic optical networks, OSA Opt. Express 27 (24) (2019) 35700–35709.

[35] T. Panayiotou, G. Ellinas, S.P. Chatzis, A data-driven QoT decision approach for multicast connections in metro optical networks, in: Proc. IEEE International Conference on Optical Network Design and Modeling (ONDM), 2016, pp. 1–6.

[36] J. Mata, I. de Miguel, R.J. Durán, J.C. Aguado, N. Merayo, L. Ruiz, P. Fernández, R.M. Lorenzo, E.J. Abril, A SVM approach for lightpath QoT estimation in optical transport networks, in: Proc. IEEE International Conference on Big Data, 2017, pp. 4795–4797.

[37] S. Aladin, A.V.S. Tran, S. Allogba, C. Tremblay, Quality of transmission estimation and short-term performance forecast of lightpaths, IEEE/OSA J. Lightwave Technol. 38 (10) (2020) 2807–2814.

[38] T. Panayiotou, H. Maryam, G. Ellinas, Deep quantile regression for QoT inference and confident decision making, in: Proc. IEEE Symposium on Computers and Communications (ISCC), 2021, pp. 1–6.

[39] R. Livni, S. Shalev-Shwartz, O. Shamir, On the computational efficiency of training neural networks, in: Proc. 27th International Conference on Neural Information Processing Systems (NIPS), 2014, pp. 855–863.

[40] J.Y. Yen, Finding the K shortest loopless paths in a network, Manage. Sci. 17 (11) (1971) 712–716.

[41] B. Shariati, A. Mastropaolo, N.P. Diamantopoulos, J.M. Rivas-Moscoso, F. Pederzolli, D. Siracusa, D. Klonidis, I. Tomkos, Spectrally-spatially flexible optical networking, in: Proc. IEEE/OSA Asia Communications and Photonics Conference, 2016, pp. 1–3.

[42] P. Ramachandran, B. Zoph, Q.V. Le, Searching for activation functions, 2017, ArXiv:1710.05941 [Cs.NE].

[43] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2017, arXiv:1412.6980 [Cs.LG].

[44] T. Panayiotou, G. Ellinas, Addressing traffic prediction uncertainty in multi-period planning optical networks, in: Optical Fiber Communication Conference (OFC), paper M3F.2, Optica Publishing Group, 2022, pp. 1–3.

Hafsa Maryam earned her M.S. degree from COMSATS University, Islamabad, Pakistan in 2017. She also served as a Visiting Lecturer at SE Department of IUSE School of Engineering & Management Sciences affiliated with UET Taxila in 2018, Rawalpindi, Pakistan. From 2018 to 2020, she has worked as a Lecturer at CS department of KICSIT sub-campus of Institute of Space and Technology (IST), and BIMS affiliated with PMAS-Arid Agriculture University, Rawalpindi, Pakistan. Currently, she is a Ph.D. student at the Department of Electrical and Computer Engineering and a researcher at KIOS Center of Excellence. Her research interest is in Intelligent Transportation System (ITS), Information Centric Networks (ICN), Content Centric Networks (CCN), Mobile Ad-hoc Networks, Wireless Communication and Fog Computing. She has authored over 10 academic articles in peer-reviewed international journals and conferences.

Tania Panayiotou received the Diploma degree in computer engineering and informatics from the University of Patras, Patras, Greece, in 2005 and the Ph.D. degree in computer engineering from the University of Cyprus (UCY), Nicosia, Cyprus, in 2013. She is currently a Research Associate with the KIOS Research and Innovation Center of Excellence, UCY. Previously, she was an Associate Researcher with the Department of Electrical Engineering, Computer Engineering and Informatics, Cyprus University of Technology, Limassol, Cyprus, and an Adjunct Lecturer with the Department of Electrical and Computer Engineering, UCY and with the Department of Information and Communication Systems, Open University of Cyprus, Latsia, Cyprus. She has authored more than 35 articles, conference papers, and book chapters. Her research interests include optical networks and transportation networks. She was the recipient of the Best Paper Award in the ONDM'17 Conference.

Georgios Ellinas holds B.Sc., M.Sc., M.Phil., and Ph.D. degrees in Electrical Engineering from Columbia University. He is a Professor at the Department of Electrical and Computer Engineering and a founding member of the KIOS Research and Innovation Center of Excellence at the University of Cyprus. Prior to joining the University of Cyprus, he also served as an Associate Professor of Electrical Engineering at City College of the City University of New York, as a Senior Network Architect at Tellium Inc., and as a Research Scientist/Senior Research Scientist in Telcordia Technologies' (formerly Bell Communications Research (Bellcore)) Optical Networking Research Group. Prof. Ellinas is a Fellow of the IET (2019), and a Senior Member of IEEE, OSA, and ACM. He has co-authored/co-edited four books on optical networks, more than 290 archived articles/conference papers/book chapters, and he is the holder of 30 patents on optical networking. His research interests are in the areas of optical networks, intelligent transportation systems, IoT, emergency response systems, and unmanned aerial systems.