

Active vs Transfer Learning Approaches for QoT Estimation with Small Training Datasets

Original

Active vs Transfer Learning Approaches for QoT Estimation with Small Training Datasets / Azzimonti, Dario; Rottondi, Cristina; Giusti, Alessandro; Tornatore, Massimo; Bianco, Andrea. - STAMPA. - (2020). (Intervento presentato al convegno OFC 2020- The Optical Networking and Communication Conference & Exhibition tenutosi a San Diego nel March 2020) [10.1364/OFC.2020.M4E.1].

Availability:

This version is available at: 11583/2826021 since: 2020-05-17T15:55:59Z

Publisher:

OSA Publishing

Published

DOI:10.1364/OFC.2020.M4E.1

Terms of use:

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

Publisher copyright

Optica Publishing Group (formely OSA) postprint/Author's Accepted Manuscript

“© 2020 Optica Publishing Group. One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modifications of the content of this paper are prohibited.”

(Article begins on next page)

Active vs Transfer Learning Approaches for QoT Estimation with Small Training Datasets

Dario Azzimonti¹, Cristina Rottondi², Alessandro Giusti¹, Massimo Tornatore³ and Andrea Bianco²

¹Dalle Molle Institute for Artificial Intelligence, Lugano, Switzerland

²Politecnico di Torino, Italy, ³Politecnico di Milano, Italy

{cristina.rottondi, andrea.bianco}@polito.it, {alessandro.giusti, dario.azzimonti}@supsi.ch, massimo.tornatore@polimi.it

Abstract: We compare the level of accuracy achieved by active learning and domain adaptation approaches for quality of transmission estimation of an unestablished lightpath, in presence of small-sized training datasets. © 2020 The Author(s)

OCIS codes: 060.4250 Fiber optics and optical communications, Networks; 060.4510 Fiber optics and optical communications, Optical communications

1. Introduction

Several supervised Machine Learning (ML) algorithms for Quality of Transmission (QoT) estimation of unestablished lightpaths have recently been investigated to overcome possible shortcomings of traditional methods such as *i*) approximated mathematical models, which usually introduce conservative design margins to compensate for uncertainties on input parameters and/or to account for simplifying modeling assumptions; *ii*) directly simulating optical signal propagation along the fiber core, which typically imposes unaffordable computational complexity if applied in real-time operational scenarios. Supervised learning algorithms for QoT estimation leverage a set of training samples, which are constituted by a vector of *features* characterizing the lightpath (e.g. length, amount of served traffic, modulation format, etc.) associated with a *target variable* that measures the QoT of the lightpath (e.g., the Bit Error Rate - BER -). It is commonly assumed that the target variable can be measured at the receiver and then collected via telemetry equipment [1]. Unfortunately these supervised algorithms require a very large set of training samples, at least in the order of hundreds or even of a few thousands of lightpaths' observations [2], to learn an accurate prediction model. However, in real installations, the collection of training samples is often hindered by practical limitations such as absence of optical monitors in some network nodes or scarcity of monitorable lightpaths in case of networks in their early deployment phase.

When the number of available samples is limited, two solutions can be applied to enlarge the training dataset. The first one consists in acquiring a limited number of additional samples via dedicated probe lightpaths (i.e., lightpaths that do not carry user traffic and are deployed with the only scope of monitoring their QoT). In this case, to maximize the accuracy improvement, we can employ *Active Learning* (AL) algorithms [5] to select route and transmission configurations of the probe lightpaths. The second one consists in acquiring additional "external" training data collected from a different network, and use *Transfer Learning* (TL) [3,4] (specifically, in this study we will use *Domain Adaptation* (DA) techniques) to extract new useful knowledge from the external data. In both cases, acquiring additional samples comes at a cost, as *i*) lightpath probes require dedicated transmission equipment and occupy spectral resources, and *ii*) acquiring external samples from another network, especially if managed by a different operator, has a cost. Hence, the choice among the two approaches is driven by both technical criteria (i.e., maximizing QoT prediction accuracy) and economic factors (i.e., the cost to collect data).

In this paper we compare the QoT estimation accuracy achieved by a few AL and DA methods on two different network topologies, as a function of the number of additional training instances that can be acquired. We describe such approaches in Sec.2 and evaluate their performance in Sec.3, adding some final remarks.

2. Transfer Learning and Domain Adaptation Approaches for QoT Estimation

In both AL and DA approaches, the features vector contains: total length of the lightpath, maximum link length, number of traversed links, amount of transmitted traffic and modulation format adopted for transmission. Moreover, we leverage a Gaussian-Processes-based classifier [5]: on input of the features of a candidate lightpath \mathbf{x}^* , the Gaussian-Process (GP) model predicts the BER value and then computes the class probabilities (where a class indicates BER value above or below a threshold γ) with the probability of excursion $p_\ell(\mathbf{x}^*) = \Phi((m_\ell(\mathbf{x}^*) - \gamma)/\sqrt{k_\ell(\mathbf{x}^*, \mathbf{x}^*)})$, where m_ℓ and k_ℓ are the posterior mean and covariance kernel of the GP model given ℓ data points and γ is the critical BER threshold. The class probabilities can be converted in

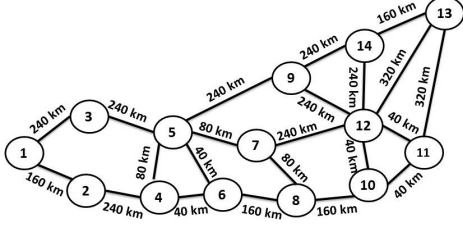


Fig. 1: Japan network topology

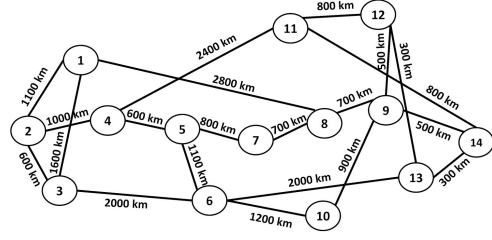


Fig. 2: NSF network topology

yes/no answers depending on whether the output probability is above or below 50% (see [5] for more details on the classifier implementation). In the following, we describe the AL and DA methods we will apply and define baselines to be used as benchmarks.

Active Learning Technique: We assume that a small training dataset \mathbf{T} containing lightpath samples from the network under consideration (referred to in the following as *target domain*) is available. Active learning aims at adaptively increasing the size of the training data by adding new samples that minimize a problem-specific acquisition function. We consider the acquisition function introduced in [7] which is related to the future variance of the classification function. A sample \mathbf{x}^{AL} that minimizes such an acquisition function also minimizes the variance of the classifier, thus leading to a better classification performance. $J_\ell(\cdot)$ denotes the acquisition function computed after ℓ samples are added to the initial training set. As in [5], we constrain the acquisition function to obtain only feasible samples. We increase the size of the training set $\mathbf{T}_0 = \mathbf{T}$ with an iterative procedure where at each step ℓ : (1) the GP-based model is trained with $\mathbf{T}_\ell = \mathbf{T}_0 \cup \{(\mathbf{x}_1^{AL}, y_1), \dots, (\mathbf{x}_\ell^{AL}, y_\ell)\}$, where y_ℓ is the BER value obtained from a probe with features \mathbf{x}_ℓ^{AL} ; (2) the model predictions are used to build the acquisition function J_ℓ ; (3) we minimize J_ℓ to select the next inputs $\mathbf{x}_{\ell+1}^{AL}$; (4) we estimate the BER value $y_{\ell+1}$ with a probe that has features $\mathbf{x}_{\ell+1}^{AL}$ and we update the training set with $\mathbf{T}_{\ell+1} = \mathbf{T}_\ell \cup \{(\mathbf{x}_{\ell+1}^{AL}, y_{\ell+1})\}$; (5) the procedure is repeated until the number of acquirable samples is exhausted.

Domain Adaptation Techniques: In the case of DA techniques, in addition to the training set \mathbf{T} , we leverage a second training set \mathbf{S} obtained from a different network (named as *source domain*). Below we briefly recall the three domain adaptation techniques considered in this work.

- *Bayesian updating* (BU) considers the model trained on the source dataset \mathbf{S} as a prior and computes a posterior distribution by updating the GP model with the data from \mathbf{T} . GP models allow for analytical updates of the posterior distribution, see, e.g., [6], thus the distribution can be updated exactly.

- *Feature Augmentation* (FA) [8] encodes the domain of a sample by tripling the length of its feature vector \mathbf{x} with a rule that depends on the domain, with the aim of exploiting both the commonalities between the two domains and the unique characteristics of each domain. More in detail, if the sample comes from \mathbf{S} , the tripled feature vector is defined as $\mathbf{x}' = \langle \mathbf{x}, \mathbf{x}, \mathbf{0} \rangle$; whereas if the sample comes from \mathbf{T} , $\mathbf{x}' = \langle \mathbf{x}, \mathbf{0}, \mathbf{x} \rangle$.

- *Correlation Alignment* (CORAL) [9] is an unsupervised technique that reduces domain shift by aligning the second-order statistics of \mathbf{S} and \mathbf{T} . The instances in \mathbf{S} are usually contained in a different manifold of the space of features than the ones in \mathbf{T} . CORAL re-aligns them by applying a transformation ϕ that re-colors the whitened features of the samples in \mathbf{S} with the covariance matrix estimated either from the features distribution of the samples in \mathbf{T} or, when available, from the samples of a set $\mathbf{T}_{\text{unlabeled}}$ of unlabeled data (i.e., samples for which the associated BER is not known) extracted from the target domain. Note that generating $\mathbf{T}_{\text{unlabeled}}$ simply requires to select the routes and transmission configuration of a large set of potential lightpaths, without measuring their BERs. The method estimates the transformation ϕ from \mathbf{S} to \mathbf{T} , then trains the classifier on $\phi(\mathbf{S}) \cup \mathbf{T}$.

Baselines approaches: We consider the following three baselines: i) *Source Domain Baseline* (SDB) trains the classifier only on \mathbf{S} ; ii) *Reduced Target Domain Baseline* (RTDB) trains the classifier only on \mathbf{T} ; iii) *Large Target Domain Baseline* (LTDB) trains the classifier on \mathbf{T}' containing a larger number of samples from the target domain.

3. Results

We consider the Japan (Jnet) and NSF (NSFnet) networks depicted in Figs. 1 and 2 and evaluate the performance of the DA and AL approaches presented in Section 2 in terms of Area Under the ROC Curve (AUC).

Synthetic dataset generation: We use the QTool presented in [2]: on input of a candidate lightpath, traffic volume and modulation format, it calculates the BER as a function of the signal-to-noise ratio measured at the channel decoder via the approximated AWGN model of dispersion uncompensated transmission over single mode fibers. The Qtool also adds random penalties to account for model uncertainties, according to an exponential distribution with average of 2 dB. We assume a flexi-grid scenario with 12.5 GHz slice width and transceivers operating at 28 Gbaud with 37.5 GHz optical bandwidth, using a modulation format chosen among dual polarization (DP)-BPSK, DP-QPSK and DP- n -QAM, with $n = 8, 16, 32, 64$. Traffic demands exceeding the capacity of a transceiver are

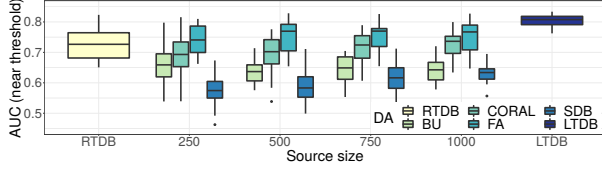


Fig. 3: Domain adaptation with target Jnet

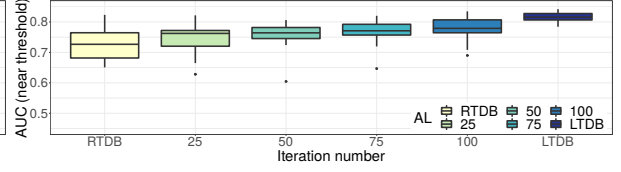


Fig. 4: Active learning with target Jnet

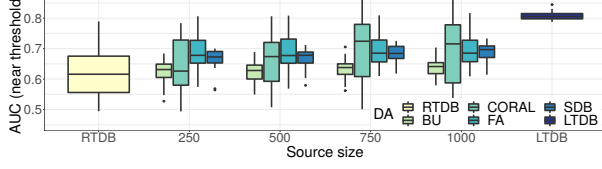


Fig. 5: Domain adaptation with target NSFnet

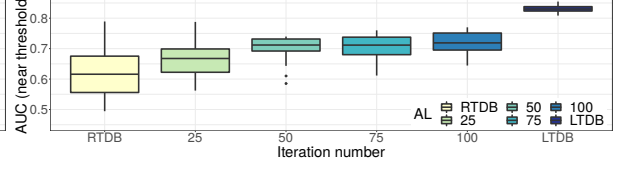


Fig. 6: Active learning with target NSFnet

served by superchannels containing multiple adjacent transceivers. For each topology, we build datasets $\mathcal{R}_{\text{source}}$ including 10000 samples generated by randomly choosing a source-destination node pair, a modulation format and a traffic demand uniformly selected in the range $[50 - 500]$ Gbps with 50Gbps granularity and computing the BER with the Qtool. We set the BER threshold to $\gamma = 4 \cdot 10^{-3}$.

Considered ML approaches: The AL and DA methods are validated against the RTDB and LTDB baselines, where LTDB assumes $|\mathbf{T}'| = 1000$. We assume $|\mathbf{S}| = 250, 500, 750, 1000$ (where $\mathbf{S} \subset \mathcal{R}_{\text{source}}$) and $|\mathbf{T}| = 25$ (where $\mathbf{T} \subset \mathcal{R}_{\text{target}}$). For CORAL only, we assume $\mathbf{T}_{\text{unlabeled}} = 1000$ (where $\mathbf{T}_{\text{unlabeled}}$ contains the feature vectors of 1000 elements in $\mathcal{R}_{\text{target}}$). For the AL algorithm, we add up to 100 new samples and we evaluate the AUC after 25, 50, 75 and 100 iterations. We evaluate the results on restricted test sets of samples from the target network with BER values in the range $[10^{-5}, 10^{-2}]$, which exhibit BER near to the threshold γ and are thus expected to be difficult to classify correctly. The test data size is 1450 (resp. 1005) for the Jnet (resp. NSFnet). Experiments are repeated 20 times, with random extraction of elements of \mathbf{S} (resp. \mathbf{T}) from set $\mathcal{R}_{\text{source}}$ (resp. $\mathcal{R}_{\text{target}}$).

Numerical assessment: Figs. 3 and 4 show the results when the target domain is the Jnet and the source domain is the NSFnet. Note that the average lightpath length in the Jnet is much smaller than in the NSFnet. SDB and BU do not adapt the source-domain data to the target domain and, in fact, their AUC does not increase when increasing the amount of data point from $|\mathbf{S}|$. Indeed, acquiring more samples of long lightpaths from NSFnet does not help classify correctly the (shorter) lightpaths of Jnet. CORAL and FA, instead, show improvements over RTDB. The AL algorithm outperforms RTDB already with 25 additional samples and achieves a median AUC of 0.78, compared to $\text{med}(\text{AUC}) = 0.73$ for RTDB and $\text{med}(\text{AUC}) = 0.82$ for LTDB. Note that FA achieves a median AUC of 0.77 with $|\mathbf{S}| = 500, 750, 1000$ and CORAL achieves $\text{med}(\text{AUC}) = 0.74$ with $|\mathbf{S}| = 1000$.

Conversely, Figs. 5 and 6 show the results with the Jnet as source domain and the NSFnet as target domain. In this case, samples coming from Jnet tend to be very useful (they help NSFnet data set to be complemented of useful information on short paths). For this reason, both BU and SDB perform relatively better here (BU, $|\mathbf{S}| = 750$, $\text{med}(\text{AUC}) = 0.64$; SDB, $|\mathbf{S}| = 750$, $\text{med}(\text{AUC}) = 0.68$) than in the previous example. CORAL ($|\mathbf{S}| = 750$, $\text{med}(\text{AUC}) = 0.72$) and FA ($|\mathbf{S}| = 750$, $\text{med}(\text{AUC}) = 0.69$) still bring improvements over RTDB ($\text{med}(\text{AUC}) = 0.62$), SDB and BU; however, the differences here are less noticeable. AL, instead, achieves a median AUC of 0.72 with 100 additional samples.

Based on the preliminary results discussed above, it emerges that comparable AUC improvements over RTDB can be obtained either with some tens of additional samples acquired from the target domain by means of an AL approach, or with some hundreds of samples gathered from another network topology. Further research will investigate the relationship between performance improvements and the degree of dissimilarity between the source and target domain (e.g. in terms of network size and transmission equipment type).

Acknowledgements: D.A. acknowledges funding from the SNSF, grant number 167199.

References

1. K. Christodoulopoulos et al., "ORCHESTRA-Optical performance monitoring enabling flexible networking," ICTON 2015.
2. C. Rottondi et al., "Machine-learning method for quality of transmission prediction of unestablished lightpaths," *IEEE/OSA JOCN*, vol. 10, no. 2, pp. A286–A297, Feb 2018.
3. W. Mo et al., "ANN-based transfer learning for QoT prediction in real-time mixed line-rate systems," *Proc. OFC* 2018.
4. L. Xia et al., "Transfer learning assisted deep neural network for OSNR estimation," *Optics Express* 27(14), pp. 19398–19406, 2019.
5. D. Azzimonti et al., "Reducing probes for quality of transmission estimation in optical networks with active learning," *Journal of Optical Communication and Networking* 12(1), pp. A38–A48, 2019.
6. C. K. Williams and C. E. Rasmussen, "Gaussian processes for machine learning," *the MIT Press*, vol. 2, no. 3, p. 4, 2006.
7. J. Bect, D. Ginsbourger, L. Li, V. Picheny, and E. Vazquez, "Sequential design of computer experiments for the estimation of a probability of failure," *Statistics and Computing*, vol. 22, no. 3, pp. 773–793, 2012.
8. H. Daumé III, "Frustratingly easy domain adaptation," *arXiv preprint arXiv:0907.1815*, 2009.
9. B. Sun et al., "Return of frustratingly easy domain adaptation," *Proc. AAAI Conf. on Artificial Intelligence* 2016.