On the benefits of domain adaptation techniques for quality of transmission estimation in optical networks

On the Benefits of Domain Adaptation Techniques for Quality of Transmission Estimation in Optical Networks

CRISTINA ROTTONDI[1,*], RICCARDO DI MARINO[1], MIRKO NAVA[2], ALESSANDRO GIUSTI[2], AND ANDREA BIANCO[1]

[1] Department of Electronics and Telecommunications, Politecnico di Torino, Turin, Italy
[2] Dalle Molle Institute for Artificial Intelligence, Lugano, Switzerland
[*] Corresponding author: cristina.rottondi@polito.com

Compiled November 11, 2020

Machine Learning (ML) is increasingly applied in optical network management, especially in cross-layer frameworks where physical layer characteristics may trigger changes at the network layer due to transmission performance measurements (Quality of Transmission - QoT) monitored by optical equipment. Leveraging ML-based QoT estimation approaches has proven to be a promising alternative to exploiting classical mathematical methods or transmission simulation tools.
However, supervised ML models rely on large representative training sets, which are often unavailable, due to the lack of the necessary telemetry equipment or of historical data. In such cases, it can be useful to use training data collected from a different network. Unfortunately, the resulting models may be uneffective when applied to the current network, if the training data (the source domain) is not well representative of the network under study (the target domain). Domain Adaptation (DA) techniques aim at tackling this issue, to make possible the transfer of knowledge among different networks.
This paper compares several DA approaches applied to the problem of estimating the QoT of an optical lightpath using a supervised machine learning approach. Results show that, when the number of samples from the target domain is limited to few dozens, DA approaches consistently outperform standard supervised ML techniques.

http://dx.doi.org/10.1364/ao.XX.XXXXXX

1. INTRODUCTION

Predicting the Quality of Transmission (QoT) of a candidate lightpath prior to its establishment plays a pivotal role for an effective design and management of optical networks. In the last few years, Machine Learning (ML) techniques for QoT estimation [1] have been advocated as promising alternatives to *i*) approximated mathematical models or *ii*) simulation frameworks that model the propagation of the optical signal along the fiber core. The former often introduce high margins to conservatively compensate for simplifying assumptions and/or for uncertainties in input parameters values, the latter typically require prohibitive computational effort when applied to real-scale scenarios. Supervised ML-based approaches learn a mapping from a set of *features*, e.g., characteristics of a lightpath such as length, amount of served traffic and adopted modulation format to a *target variable*, e.g., an indicator of the expected QoT in the lightpath, such as the Bit Error Rate (BER). To make the learning phase effective, a large amount of samples (training set) must be provided to the learning algorithm. Each sample consists of the features of an already established lightpath associated with the actual value of the target variable (*ground truth*), e.g. the BER that can be measured at the receiver node [2].

On one hand, supervised ML methods assume that the training set is large and *representative* of the samples that will be processed when the model is exploited to predict the QoT. On the other hand, the collection of training samples is often hindered by practical issues (e.g., lack of dedicated telemetry equipment in every network node) or is too costly to permit the acquisition of large datasets. This is especially true for networks in the early stage of deployment, where the number of already installed (and thus, monitorable) lightpaths is very limited. Some approaches to cope with the scarcity of training data have recently been proposed. The usage of synthetically generated data has been suggested for the initial training phase of ML algorithms, which will be re-trained once field-gathered data become available to mitigate the model inaccuracies introduced by the artificial

data generation process [3]. Alternatively, active learning may be leveraged to complement the data at disposal with a small amount of samples acquired by means of dedicated probes to cover specific areas of the feature space, indicated by the active learning model itself [4]. However, it is sometimes possible to rely on large training datasets from a different network than the one on which the ML model should operate. Therefore, we assume that a large amount of training data is given for a *source domain* (e.g., a backbone network monitored for a long operational period), and must be used to train a model that predicts the QoT of lightpaths to be established in a different *target domain* (e.g., a newly deployed network), for which only a small labeled training dataset is available [5–7]. In such a scenario, we wish to exploit at best the data from the source domain to tailor a good model to the target domain. This approach is known in ML research as *Domain Adaptation* (DA) [8] and leverages the intuition that samples from one domain provide useful information concerning the solution of the QoT estimation problem in the other domain.

Intuitively, the more the target domain differs from the source domain, in terms of distribution of lightpath lengths and/or in the type of installed fibers and transmission equipment, the less accurate the prediction on the lightpaths drawn from the target domain is expected to be. Indeed, the mapping between features and target variable somewhat differs between the two domains: the joint probability distribution of features and target variable is not the same in the two domains. Therefore, training data extracted from the source domain need to be complemented by a few samples obtained from the target domain: quantifying the amount of samples required to achieve satisfactory predictive performance of the adopted learning model is a crucial issue to determine the practical applicability of DA techniques in real-world scenarios.

In [9], a preliminary assessment of existing DA approaches for ML-based QoT estimation of candidate lightpaths was presented. The focus was on two networks characterized by different topologies, but adopting the same fiber type and transmission equipment. In this paper, we generalize that analysis by considering various topologies, rescaled in their link length by different factors, to obtain a wider set of domains. We assess the performance of two DA techniques as a function of the number of available training instances from the target domain and on the degree of dissimilarity between the two domains. Furthermore, we compare the performance of different learning algorithms to identify which benefits more of the application of DA techniques.

The remainder of the manuscript is organized as follows: after a brief review of the related literature in Section 2, we describe the applied DA methods in Section 3 and present the adopted classification framework in Section 4. We assess the performance of the considered approaches in Section 5 and finally draw our conclusions in Section 6.

2. RELATED WORK

Though ML approaches for QoT estimation have been widely investigated in the last few years [1], only a few studies proposing the use of Transfer Learning (TL) to predict various metrics related to QoT (e.g., BER, Q-factor or Optical Signal to Noise Ratio - OSNR) or network design/management (e.g., routing reconfiguration time) have appeared.

In [5], an Artificial Neural Network (ANN) based TL framework is adopted to predict the Q-factor in a real-time mixed line-rate optical system. A 4-span Large Effective Area Fiber (LEAF) with Quadrature Phase Shift Keying (QPSK) transmission is considered as source domain, while a 4-span LEAF with 16-Quadrature Amplitude Modulation (QAM) transmission, a 2-span LEAF with 16-QAM transmission and a 3-span Dispersion-Shifted Fiber (DSF) with QPSK transmission are used as target domains. After an initial training phase during which ANN weights are learned using samples from the source domain, the pre-trained weights of the hidden layer(s) are readjusted based on a small number of samples from one of the three different target domains, thus tailoring the trained model to the target domain without need to recompute the weights from scratch. Results show that only a few dozens of training samples are necessary to fine-tune the weights in order to produce an accurate prediction of QoT for the three domains, applying knowledge transfer.

In [6], a Deep Neural Network (DNN) for OSNR estimation is proposed, which relies on TL techniques to enable fast remodeling in case of variations of different system parameters such as optical launch power, chromatic dispersion and bit rate. Once weights and bias values are obtained from the initial training phase, they are readjusted leveraging other samples obtained by changing launch power, residual dispersion and bit rate. Also in this case, the TL mechanism exploits the reuse of ANN parameters obtained from a pre-trained model in a new model. Results show that a consistent reduction in the number of required samples from the target domain and in the training time can be achieved, with respect to training the ANN from scratch with randomly initialized weights scheme, without relying on previously acquired knowledge from the source domain.

Using a similar methodological approach, in [10] a multi-task DNN is adopted for joint OSNR prediction and modulation format identification. Results obtained from an experimental test-bed show that knowledge transfer of the neural network weights achieves a significant reduction of the required training samples and epochs.

In [11, 12] a TL methodology for QoT estimation in multi-domain networks with broker orchestration is discussed. QoT monitoring data of intra-domain paths collected by each network domain manager are associated to a set of encoded features (to avoid disclosure of confidential domain information) and provided as samples to the broker plane. The broker leverages the samples acquired from various domains to train a DNN-based QoT estimator for each inter-domain path. In this context, the transfer of trained weights between neural networks from a model trained with samples from a different path is limited to the first few hidden layers and those weights are kept locked in the new model. In the upper hidden layers, conversely, weights are randomly initialized and trainable with data from the target domain. Note that in [12] the architecture of the DNN is defined beforehand by applying a genetic algorithm which optimizes the number of layers, the number of neurons per layer, the activation and learning functions, with an evolutionary approach. Results show that the number of required training samples can be reduced by one order of magnitude.

In [7], a resource reservation algorithm for traffic requests in Space Division Multiplexed Elastic Optical Networks (SDM-EON) based on transductive TL is proposed. The study considers a scenario where traffic requests either require immediate reservation or can be reserved in advance. If the latter type of requests cannot be accommodated, TL is used to predict the spectrum de-fragmentation time to complete resource reallocation before their start time. A set of traffic requests is collected

from the source domain topology and then sorted to discriminate between *i*) blocked service set, consisting of all requests that will be blocked and not be accounted for the reservation of the network resources; *ii*) affected service set, including all the requests already satisfied in the network. Each element of the sets is associated to a value indicating the duration of spectrum de-fragmentation process. These samples are used to perform the initial training phase. Then, the blocked service set and affected service set from the selected target domain topology are used as inputs to refine the pre-trained model. The outputs of the model coincide with the prediction of spectrum migration time of the requests. Finally, based on this output, a further optimization phase of spectrum utilization takes place to improve resource reservations among different services. Authors used a six-node topology as source domain, while the fourteen-node NSF topology is used as target domain. They show that the adoption of TL may decrease the blocking probability up to 67%.

TL methodologies have also been applied to other management tasks in optical networks: in [13], TL-aided neural networks are proposed for nonlinear equalization in a short optical link. The weights obtained by training the neural network with samples acquired from a source system with different bit rate or link length are leveraged to initialize the training phase of the equalizer to be applied in the target system, thus allowing for a drastic reduction in the number of training epochs and required samples.

All the studies mentioned so far, including ours, apply knowledge transfer approaches in a supervised learning framework. In the context of reinforcement learning, a multi-task-learning-aided knowledge transfer approach is proposed in [14] to learn routing, modulation format and spectrum assignment policies: given the source-destination nodes and the time duration of a traffic request, the agent learns how to choose the best route among a set of pre-computed paths and performs first-fit spectrum assignment. Results show that knowledge acquired from paths belonging to a given network topology can be leveraged to take routing, modulation and spectrum assignment decisions for traffic requests to be allocated in a different network topology, achieving up to 20% reduction of the blocking probability.

In this study, we tackle the problem of BER classification, i.e., predicting whether the BER of an un-established lightpath will fall below or above a given system threshold. In our analysis we will consider various network topologies, i.e. the NSF and Japan networks, which we opportunely rescale to account for populations of lightpaths of different lengths. Differently from the previously mentioned studies, our features are high-level lightpath characteristics such as length, amount of served traffic and modulation format adopted for transmission. The TL approaches mentioned above rely on the standard technique of fine-tuning pre-trained neural network models on the target domain. In contrast, this paper explores DA techniques that explicitly account for diverse feature distributions from different domains.

A recent trend in ML deals with models [15, 16] that explicitly operate on graph data; when adopting these models, the topology of different optical networks could be considered an additional input to the model rather than a domain. By training such a model on data from many different optical networks of known topology, this approach could yield a model that generalizes to new unseen networks. In this paper we do not follow this approach and rely on standard classifiers designed to operate on low-dimensional feature vectors; when a new network is introduced, the model needs to be retrained (with domain adaptation) using ad-hoc data collected from the new network. As in our case the instances are paths within a given network, instead of handling the entire network as an input, which would severely limit the size of the training set, we handle the problem as a domain adaptation task.

3. DOMAIN ADAPTATION APPROACHES FOR QOT ESTIMATION

We describe the DA techniques adopted in this study and compare them to simple ML algorithms that will be considered as benchmarks for their performance evaluation.

A. Baselines

Fig.1a-c depicts the three scenarios against which the adopted DA techniques will be benchmarked. In the following, we consider a large set $\mathbf{S}$ including samples collected from the source domain $\mathcal{R}_{\text{source}}$ that is used to train a QoT estimation model and a small set $\mathbf{T}$ of samples gathered from the target domain $\mathcal{R}_{\text{target}}$. We consider a generic supervised learning framework and define three benchmark scenarios:

- **Only Source Domain Baseline** (SDB) trains the model only on $\mathbf{S}$ (Fig.1a).
- **Only Target Domain Baseline** (TDB) trains the model only on $\mathbf{T}$ (Fig.1b).
- **Dataset Mixing Baseline** (DMB) trains the classifier on $\mathbf{S} \cup \mathbf{T}$ (Fig.1c).

For validation purposes, in all scenarios the test phase is conducted using a large set $\mathbf{T}'$ gathered from the target domain $\mathcal{R}_{\text{target}}$, containing samples that were not used during the training phase (i.e., $\mathbf{T} \cap \mathbf{T}' = \varnothing$).

We expect SDB, even if trained on a large amount of data, to perform badly on $\mathbf{T}'$ because of domain shift. TDB is not affected by domain shift, but suffers the limited size of the training set that might cause overfitting. DMB should improve over SDB, but the effect of $\mathbf{T}$ will be "washed out" by the large size of $\mathbf{S}$.

B. Domain Adaptation Techniques

The two DA techniques adopted in this study are Feature Augmentation (FA) [17] and CORrelation ALignment (CORAL) [18].

FA (see Fig.1d) implements a simple approach which encodes the domain of a sample by augmenting its feature vector. In particular, the length of the original feature vector $\mathbf{x}$ is tripled, with a rule that depends on the domain: if the sample belongs to $\mathbf{S}$, the resulting feature vector is computed as $\mathbf{x}' = \langle \mathbf{x}, \mathbf{x}, \mathbf{0} \rangle$; otherwise, for a sample in $\mathbf{T}$, the feature vector is redefined as $\mathbf{x}' = \langle \mathbf{x}, \mathbf{0}, \mathbf{x} \rangle$. This augmentation transformation is applied to all samples, both in the training and test phases. It is expected that such transformation allows a classifier to learn and exploit both the commonalities between the two domains and the unique characteristics of each domain.

CORAL is originally conceived as an unsupervised DA technique that minimizes domain shift by aligning the second-order statistics of source and target data-sets. This is done by applying a transformation ϕ that re-colors whitened source features of the source domain with the covariance of the distribution of the dataset gathered from the target domain. In this case, we assume that a large amount $\mathbf{T}_{\text{unlabeled}}$ of unlabeled data (i.e., samples for which the associated BER is not known) from the target domain is available, and we implement the following steps (see Fig.1e):

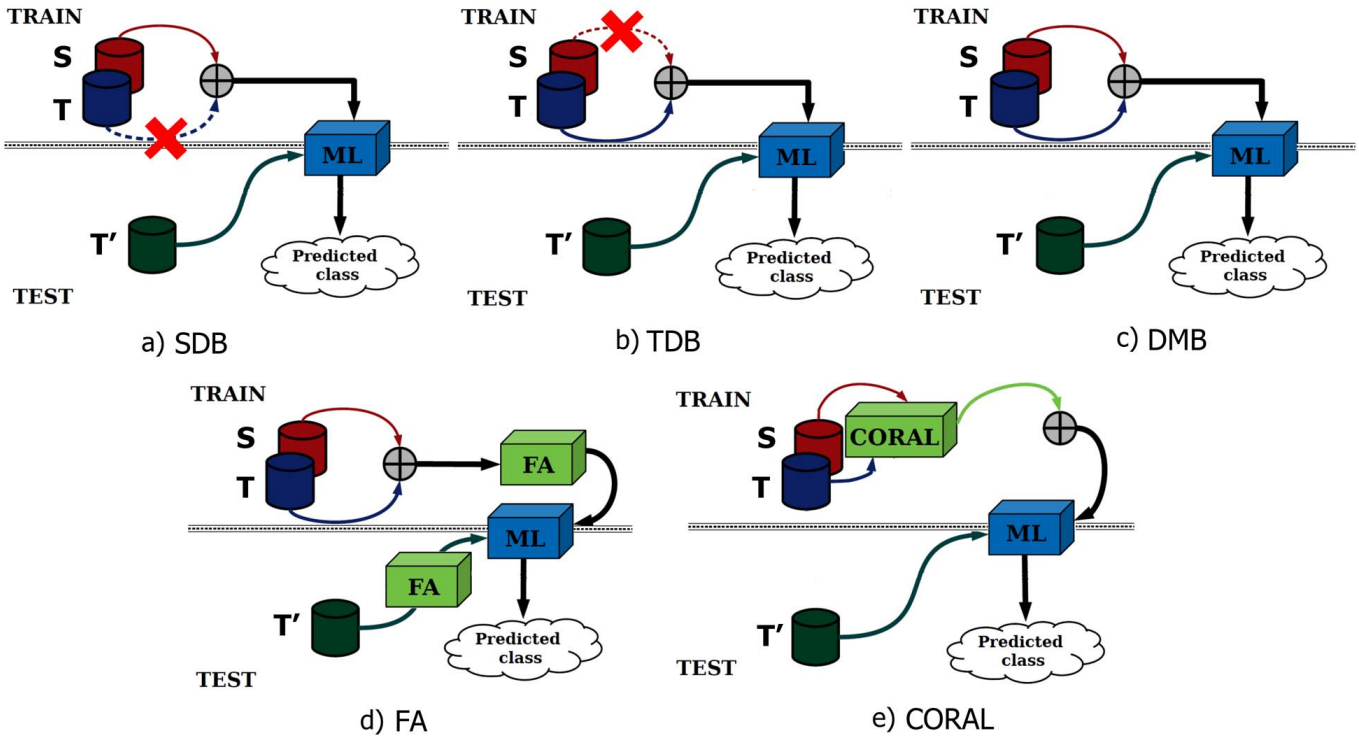

Fig. 1. Baseline and Domain Adaptation learning scenarios.

1. estimate the transformation ϕ from the source to the target feature spaces using $\mathbf{S}$ and $\mathbf{T}_{\text{unlabeled}}$;
2. train the ML algorithm on $\phi(\mathbf{S})$.

Note that, differently from FA, CORAL can be applied even when no BER measurements from the target domain are available, because large sets of prospective lightpath configurations can be obtained without proceeding with their establishment in the network.

As in the former three cases, validation is performed using a large set $\mathbf{T}'$ gathered from $\mathcal{R}_{\text{source}}$, disjoint from $\mathbf{T}$.

4. CLASSIFICATION FRAMEWORK AND SYNTHETIC DATA GENERATION TOOL

A. ML-Classifier

We adopt the ML framework for QoT classification proposed in [19], which uses the following features to characterize a lightpath: total length, number of traversed links, maximum link length, amount of traffic to be transmitted and modulation format to be adopted for transmission. Optionally, the following additional features can be included to characterize the lightpaths' neighbor channels: traffic volume, modulation format and guardband size of the spectrally-nearest right and left adjacent channels co-propagating along at least one of the links traversed by the considered lightpath. Given a candidate lightpath, the classifier provides as output a binary variable whose value depends on whether the predicted probability that the lightpath configuration satisfies a given threshold γ on the BER measured at the receiver is above or below 50%. In this study we use Support Vector Machine (SVM) classifiers, which proved to benefit from the application of DA techniques more substantially than Random Forests (RF), adopted in [19]. In Section 5D we also report results comparing different models.

B. E-Tool Data Generator

To generate synthetic BER measurements we use the E-Tool presented in [19]. Given a candidate lightpath, traffic volume and modulation format, the E-Tool calculates the BER as a function of the signal-to-noise ratio measured at the channel decoder via the approximated additive white Gaussian noise model of dispersion uncompensated transmission over single mode fibers. The E-Tool also adds random penalties (with negative exponential distribution and average 1 dB) to account for model uncertainties. 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, QPSK and n-QAM, with $n = 8, 16, 32, 64$. Traffic demands exceeding the capacity of a transceiver are served by super-channels containing multiple adjacent transceivers.

C. Dataset generation

We consider the Japan (Jnet) and NSF (NSFnet) networks depicted in Fig.2. In our numerical assessment, each topology has been resized by multiplying the length of each link by a scaling factor $\alpha = 0.25, 0.5, 1, 2$, thus obtaining multiple rescaled copies of the two network graphs.

To construct the training dataset for each topology, we produce instances by randomly choosing a source-destination node pair, a modulation format and a traffic demand uniformly selected in the range $[50 - 500]$ Gbps with 50 Gbps granularity and calculating the BER with the E-Tool. The choice of a 50 Gbps granularity is based on the transmission capacity of a 28 GBaud transceiver operating with DP-BPSK modulation format, i.e., the basic transceiver configuration adopting the least spectrally efficient modulation format. We set the BER threshold to $\gamma = 4 \cdot 10^{-3}$. The distributions of lightpath lengths obtained in each topology are plotted in Fig.3, for different values of the

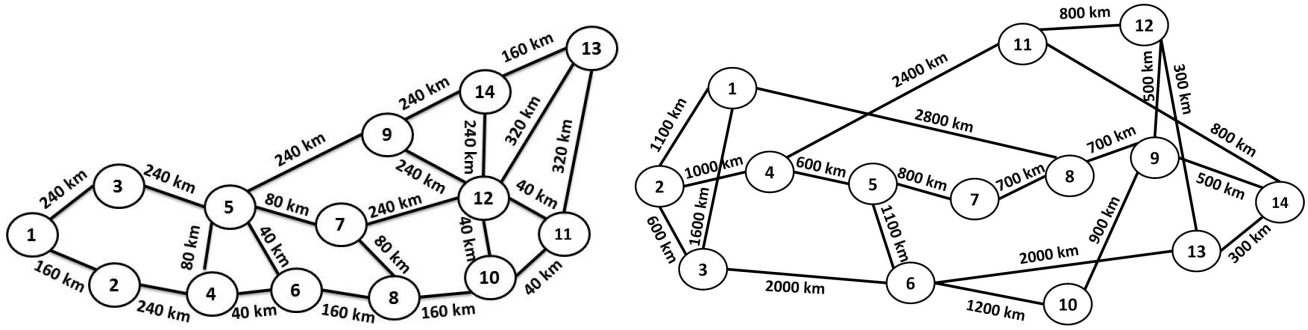

Fig. 2. Japan (left) and NSF (right) network topologies

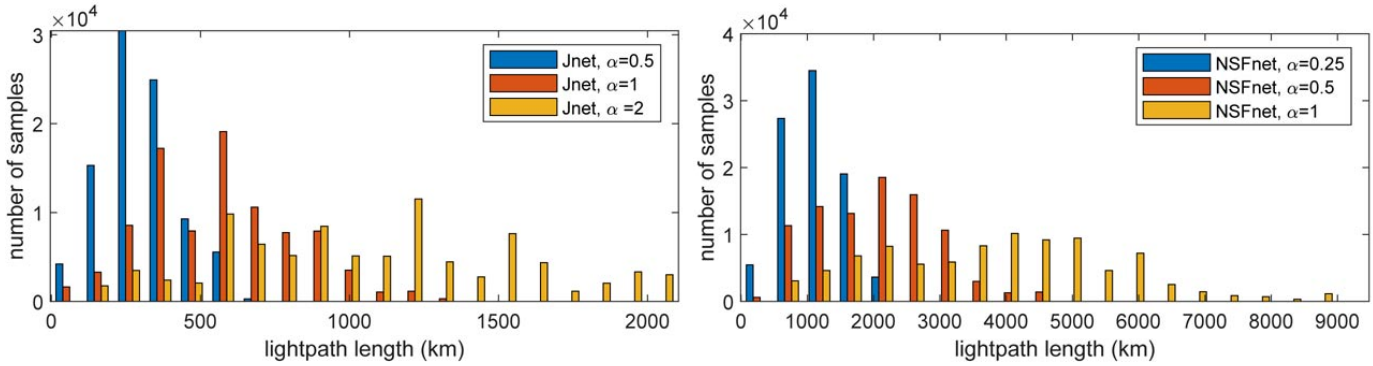

Fig. 3. Lightpath length distributions in the rescaled Jnet topologies (left) and NSFnet topologies (right), considering 90000 samples

scaling factor α. Note that the values of the α parameter were chosen with the aim of covering a variety of ligthpath lengths distributions in the two domains. In particular, the domain pair constituted by the Jnet with scaling factor $\alpha = 2$ and the NSFnet with scaling factor $\alpha = 0.25$ show very similar ligthpath lengths distributions, both in the range $[1 - 2100]$ km, whereas the pair constituted by the Jnet scaled by $\alpha = 0.5$ and the NSFnet scaled by $\alpha = 1$ offers the most diverse lightpath length distributions, with Jnet lightpaths spanning up to 700 km and the NSFnet lightpaths ranging up to 9000 km (i.e., lengths differ on average by one order of magnitude).

The test dataset $\mathbf{T}'$ is always constructed, for both topologies, by generating a separate set of instances $|\mathbf{T}'| = 45000$ with the same procedure used to produce the training dataset.

D. Considered scenarios

The two DA methods, FA and CORAL, are benchmarked against the previously described SDB, TDB and DMB baselines and against a Large TDB (LTDB) baseline that trains the SVM classifier using $|\mathbf{T}| = 10000$ samples drawn from the target domain $\mathcal{R}_{\text{target}}$. The SDB baseline and the two DA methods assume $|\mathbf{S}| = 10000$ for the training phase, where the samples of $\mathbf{S}$ are extracted from $\mathcal{R}_{\text{source}}$. In the TDB and DMB baselines and in the FA approach, $|\mathbf{T}| = 10, 50, 100, 500$, where the samples of $\mathbf{T}$ belong to the domain $\mathcal{R}_{\text{target}}$. For CORAL only, we assume $|\mathbf{T}_{\text{unlabeled}}| = 10000$, where $\mathbf{T}_{\text{unlabeled}}$ contains the feature vectors of elements belonging to $\mathcal{R}_{\text{target}}$. Indeed, for the QoT estimation task, collecting unlabeled samples of the set $\mathbf{T}_{\text{unlabeled}}$ from the target domain is trivial, as we simply need to select route, traffic volume and modulation format of a perspective lightpath to derive its feature vector, without measuring its BER. If not differently stated, we include in the feature vector the lightpath total length, number of traversed links, maximum link length, amount of traffic to be transmitted and modulation format to be adopted for transmission. Experiments are repeated 20 times, with random extraction of the elements of $\mathbf{T}$ from the pool of training samples belonging to domain $\mathcal{R}_{\text{target}}$.

5. NUMERICAL ASSESSMENT

A. Performance of Feature Augmentation

We start analyzing the performance of the FA approach. Fig.4 plots on the left side the Area Under the ROC Curve (AUC, where ROC indicates the Receiver Operating Characteristic) obtained when the original Jnet ($\alpha = 1$) is used as source domain and the NSFnet rescaled by a factor $\alpha = 0.25$ as target domain. As noticeable from Fig.3, the length of the lightpaths deployed in the Jnet with $\alpha = 1$ spans the range up to 1300 km, whereas lightpaths deployed in the NSFnet with $\alpha = 0.25$ exhibit lengths up to 2000 km. Since the average lightpath length in the Jnet is lower than in the NSF network, little knowledge about the BER of long lightpaths can be obtained through the samples of the Jnet, i.e., the source domain used to compute SDB. This motivates the AUC gap between SDB and LTDB reported in Fig.4 and for this reason, when $|\mathbf{T}| \geq 50$, TDB always outperforms SDB, hinting that learning even from a limited number of samples gathered from the target domain yields to better knowledge about long lightpaths rather then relying on a large amount of samples of short lightpaths obtained from the source domain. Indeed, considering the performance of DMB, it turns out that learning from $\mathbf{S} \cup \mathbf{T}$ leads to worse AUC values with respect to learning exclusively from $\mathbf{T}$. Conversely, FA outperforms SDB and DMB when $|\mathbf{T}| \leq 50$, whereas for large $|\mathbf{T}|$, FA shows comparable performance to TDB and closely approach LTDB, hinting that 100 samples gathered from the target domain are sufficient to obtain classification results almost as accurate as

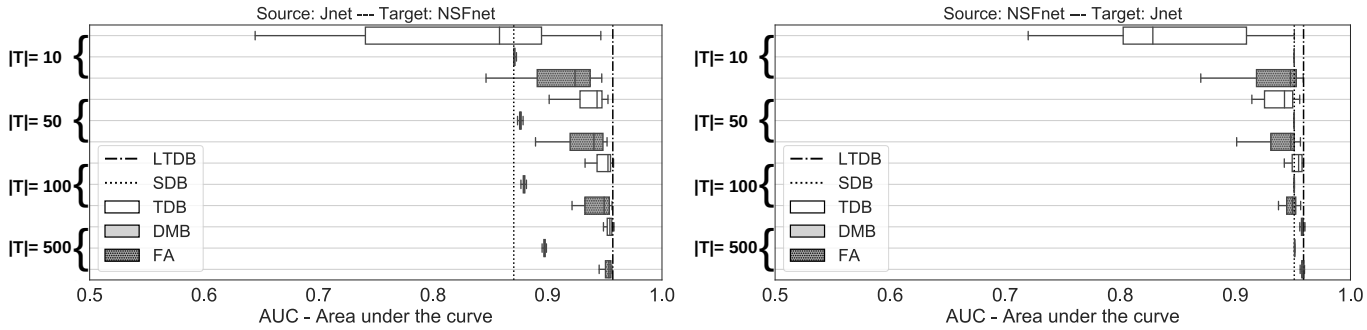

Fig. 4. AUC comparison using Jnet ($\alpha = 1$) as source domain and NSFnet ($\alpha = 0.25$) as target domain dataset (left), and viceversa (right).

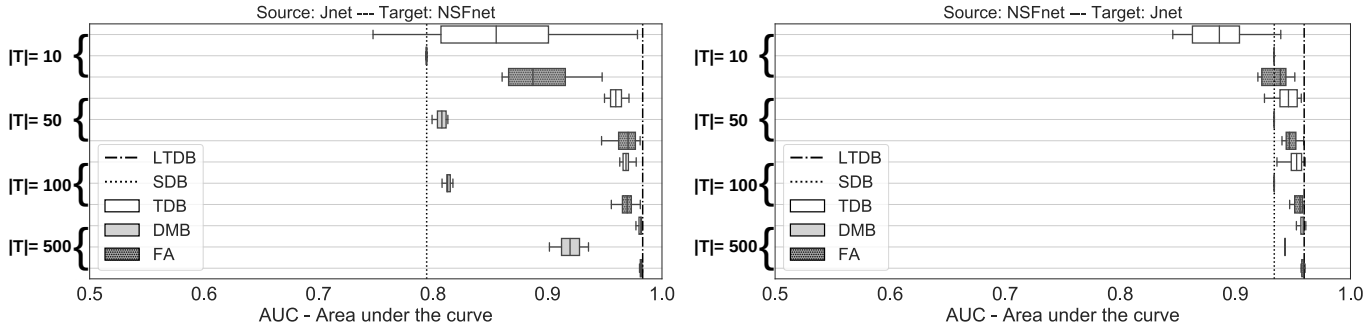

Fig. 5. AUC comparison using Jnet ($\alpha = 1$) as source domain and NSFnet ($\alpha = 1$) as target domain dataset (left), and viceversa (right).

those obtained by LTDB with 10000 samples.

In Fig.4, we report on the right side AUC results obtained using the NSFnet with $\alpha = 0.25$ as source domain and the Jnet with $\alpha = 1$ as target domain. As the average lightpath length in the NSF network is higher than in the Japan network, set $|\mathbf{S}|$ is expected to contain a representative number of lightpath samples for the whole range of lengths that can appear in the Jnet. For this reason, the gap between the AUC values obtained by SDB and LTDB baselines is very small (around 0.01), meaning that learning from samples gathered from the NSFnet already provides good knowledge on the BER of lightpaths deployed in the Jnet. Moreover, when $|\mathbf{T}|$ is low, SDB outperforms TDB, confirming that learning from a high number of source domain samples is more effective than learning from a very limited number of target domain samples. Additionally, learning from the set $\mathbf{S} \cup \mathbf{T}$ (DMB) brings no improvements in the AUC with respect to SDB. In this context, when $|\mathbf{T}|$ is low, FA significantly improves the AUC with respect to TDB, but does not reach the values exhibited by SDB and DMB baselines. Conversely, when $|\mathbf{T}| = 500$, FA and TDB provide similar results, which closely approach LTDB. It follows that, in such a scenario, DA techniques are expected to be less useful, as samples gathered from the source domain are already sufficiently representative of the part of the feature space which is relevant for the target domain.

We now repeat the same analysis focusing on the original Jnet and NSFnet topologies (i.e., $\alpha = 1$ for both networks) and report the obtained results in Fig.5. In this case, the lightpath lengths in the two topologies cover two significantly different ranges, since the NSFnet exhibits lightpaths up to 8000 km long (see right side of Fig.2), but only a few of them cover the range below 1300 km, where the Jnet lightpaths lie (see left side of Fig.2). Therefore, in both cases the AUC gaps between SDB and LTDB are significant (especially when the Jnet is used as source domain) and when $|\mathbf{T}|$ is low the benefits of adopting FA emerge more clearly, if compared to DMB and TDB. Moreover, in Fig.6 we consider the same scenarios of Fig.5, but we restrict the test set to 3000 elements exhibiting a BER within the range $[10^{-5}, 10^{-2}]$, i.e., close to the system threshold γ. Such samples are expected to be more difficult to classify correctly. Though the AUC is generally lower, the considered approaches show AUC trends similar to those obtained considering the full test dataset and in this case FA shows some improvements over TDB also for large cardinalities of $\mathbf{T}$.

Finally, to more thoroughly evaluate the impact of the difference between source and domain network sizes on the performance of FA, in Fig.7 we report on the left side the performance of FA, DMB, SDB, TDB and LTDB using Jnet ($\alpha = 0.5, 1, 2$) as source domain and NSFnet ($\alpha = 1$) as target domain dataset, assuming that $|\mathbf{T}| = 10$, whereas on the right side we report the same results considering NSFnet ($\alpha = 0.25, 0.5, 1$) as source domain and Jnet ($\alpha = 1$) as target domain. In both scenarios, decreasing the intersection between the supports of lightpath length distributions yield increasing gaps between the performance of SDB and LTDB. DMB never improves over SDB. When Jnet is used as source domain, FA achieves better AUC than TDB as long as the supports of the two distributions have a reasonably large intersection, but when the size of Jnet is halved (i.e., when the length of its links is multiplied by $\alpha = 0.5$) the median AUC is lower in FA than in SDB. Conversely, when the NSFnet is used as source domain, FA shows improvements over TDB for all scaling factors and it achieves higher AUC than SDB when

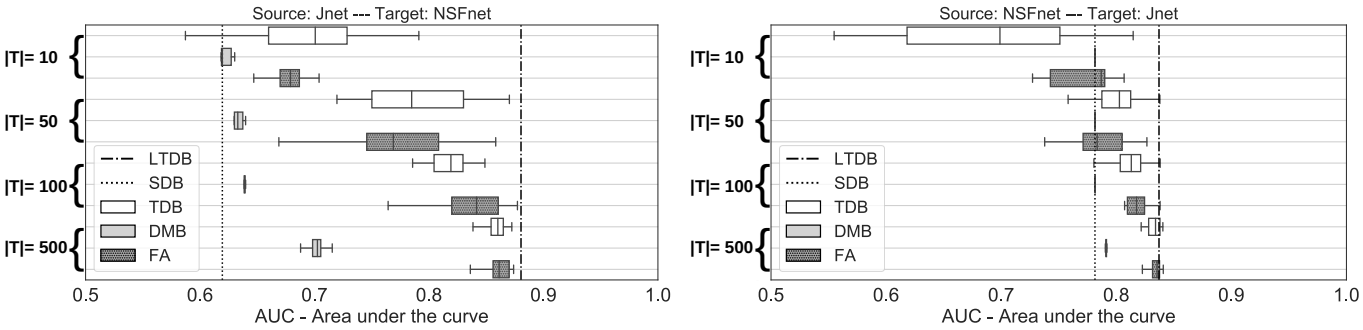

Fig. 6. AUC comparison over near-to-threshold test samples using Jnet ($\alpha = 1$) as source domain and NSFnet ($\alpha = 1$) as target domain dataset (left), and viceversa (right).

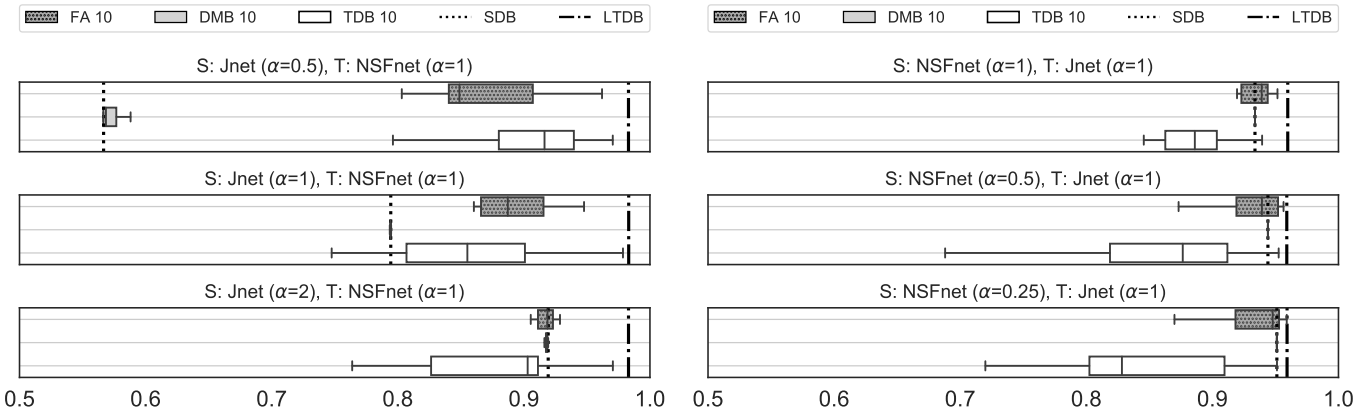

Fig. 7. AUC comparison with $|\mathbf{T}| = 10$, using Jnet ($\alpha = 0.5, 1, 2$) as source domain and NSFnet ($\alpha = 1$) as target domain dataset (left), or using NSFnet ($\alpha = 1, 0.5, 0.25$) as source domain and Jnet ($\alpha = 1$) as target domain dataset (right).

the NSFnet is much larger than the Jnet (i.e., when the length of the NSFnet links is multiplied by $\alpha = 1$).

B. Performance of Correlation Alignment

We now focus on the CORAL approach and compare its performance to that of SDB and LTDB. Note that, since CORAL relies exclusively on sets $\mathbf{S}$ and $\mathbf{T}_{\text{unlabeled}}$, both having fixed size in our experiments, results are independent of the cardinality of $\mathbf{T}$. Fig.8 shows on the left side the AUC obtained by CORAL when the Jnet (with different scaling factors) is used as source domain and the NSFnet (with $\alpha = 1$) is used as target domain, whereas on the right side the NSFnet (with different scaling factors) is used as source domain and the Jnet (with $\alpha = 1$) is used as target domain. When Jnet is used as source domain, CORAL always closely approaches the LTDB baseline, whereas the AUC values achieved by SDB decrease when the overlap of the lightpath length distributions in source and target domain diminishes. Therefore, CORAL proves to be useful especially in conditions where the lightpath lengths of the source domain are significantly lower than those of the target domain. On the right side, where results obtained by swapping source and target domain topologies are shown, the AUC achieved by SDB is quite close to that of LTDB for all values of α; the limited gap leaves very little opportunity for CORAL to yield improvements over the SDB baseline.

Based on the above reported results, we conclude that both DA approaches show considerable improvement in the AUC with respect to standard ML techniques, when the number of available samples from the target domain is very limited and when the source domain network has smaller size than the target domain network.

C. Impact of the features choice

In this subsection we evaluate the effect of enlarging the feature vector to include characteristics of the neighboring channels of the lightpath being considered. Therefore, we repeat the same experiments described in Subsections 5A and 5B including as attributes the traffic volume, modulation format and guardband size of the spectrally-nearest right and left adjacent channels co-propagating along at least one of the links traversed by the considered lightpath. Results obtained with FA using the Jnet with $\alpha = 1$ as source domain and the NSFnet with $\alpha = 1$ as target domain, and viceversa, are reported in Fig.9. The general trend is very similar to that shown in Fig.5, though the AUC values obtained by TDB and FA are slightly lower when the number of samples is 100 or less. This is due to the fact that with the increase in the number of dimensions of the feature space, the classification problem becomes more complex and more samples are likely to be required, in order to achieve comparable AUC performances. When we focus on samples near to the system threshold γ, comparing results reported in Fig.10 to those shown in Fig.6 we observe that, by using 11 features instead of 5, the AUC values achieved by TBD with $|T| = 10$ drop significantly (median value below 0.6 instead of around 0.7), whereas FA

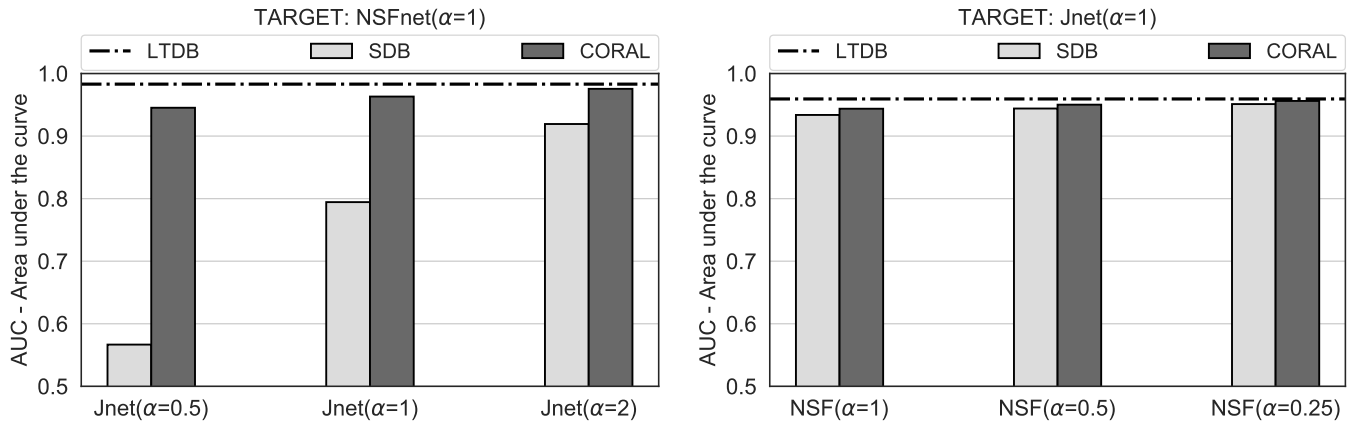

Fig. 8. AUC performance of CORAL.

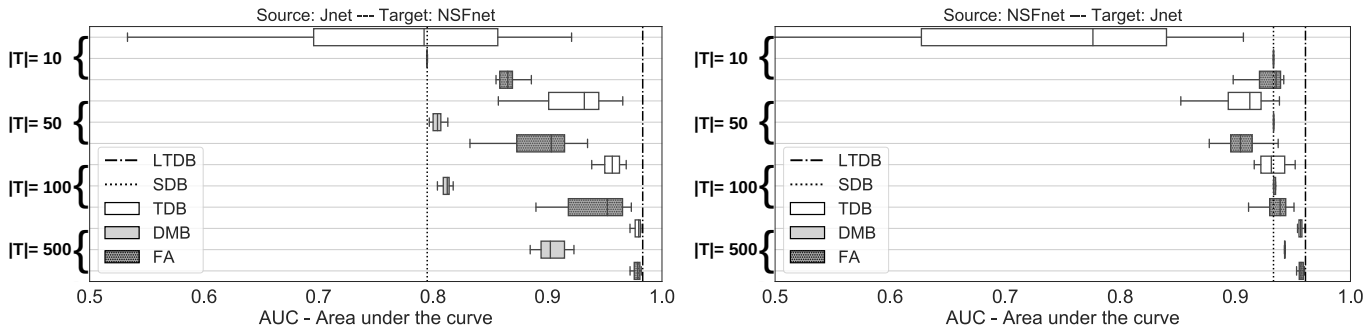

Fig. 9. AUC comparison using Jnet ($\alpha = 1$) as source domain and NSFnet ($\alpha = 1$) as target domain dataset (left), and viceversa (right), when features characterizing neighbor channels are included.

maintains similar AUC values (of at least 0.7) in both cases.

Results obtained with FA using 11 features in all the remaining scenarios showed analogous characteristics to those discussed above and are thus not reported for the sake of conciseness. The same holds for the results obtained with CORAL using 11 features (also not reported), which do not show significant differences with respect to those appearing in Fig.8.

D. Impact of the learning model

To conclude our analysis, we consider two alternative estimators, i.e., Logistic Regression (LR) and RF classifiers, in place of the SVM adopted in the previous experiments. This comparison is possible because the considered DA techniques are agnostic to the downstream classifier. We repeat the experiments for the three baselines and the two DA techniques considered so far. In Fig.11 we report results obtained in the same scenario used for the SVM-based classifier in Fig.5(left). While LR achieves comparable results to those obtained with SVM, RF shows a very different trend, with SDB and LTDB providing exactly the same AUC value (0.97). In such case, there is no benefit in applying DA techniques, since it is possible to obtain a very accurate estimate relying exclusively on the samples gathered from the source domain. In fact, RF is a powerful estimator, which is capable of representing the training set more accurately than LR. Furthermore, as an ensemble method, RF resists overfitting, which in this case helps its generalization ability even on a different domain. In this context, the effect of DA approaches is beneficial to LR but irrelevant for RF.

6. CONCLUSION

In this paper, we considered a ML-based QoT estimation frameworks for candidate optical paths, focusing on the particular case where large training datasets are unavailable or their acquisition requires high costs. In such scenarios, it is necessary to exploit the few available data and integrate them with larger datasets gathered from a different network domain. To this aim, DA techniques can be adopted to transform the features space based on the domain from which the data have been collected. We assessed the performance of two types of DA approaches in different settings, depending on the number of available samples belonging to the target domain. Achieved results show that Feature Augmentation and Correlation Alignment outperform a model trained by simply joining the sets of data collected from source and target domains. Furthermore, we explored the dependency of the AUC on the overlap between the two distributions of lightpath lengths of the two domains. Results show that DA techniques reduce the AUC decrement when the overlap between the two distributions becomes smaller.

FUNDING

This work is partially supported by the Swiss National Science Foundation (SNSF) through the NCCR Robotics.

REFERENCES

1. F. Musumeci, C. Rottondi, A. Nag, I. Macaluso, D. Zibar, M. Ruffini, and M. Tornatore, “An overview on application of machine learning

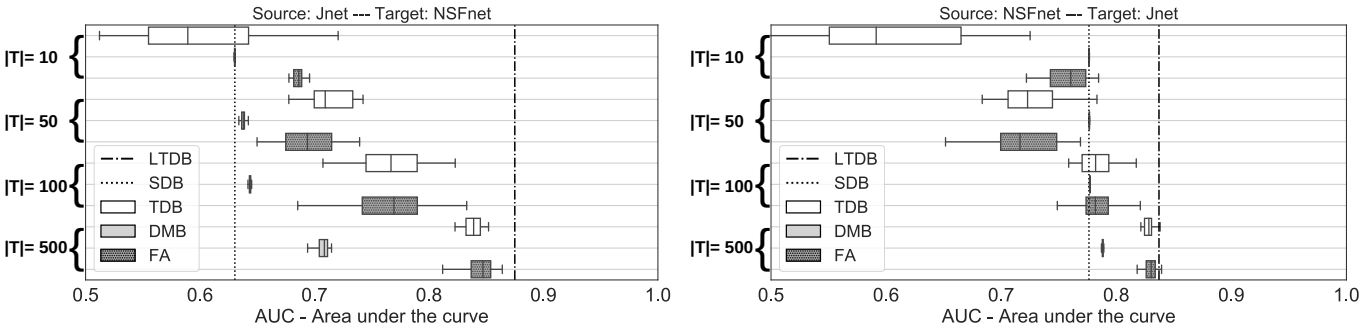

Fig. 10. AUC comparison over near-to-threshold test samples using Jnet ($\alpha = 1$) as source domain and NSFnet ($\alpha = 1$) as target domain dataset (left), and viceversa (right), when features characterizing neighbor channels are included.

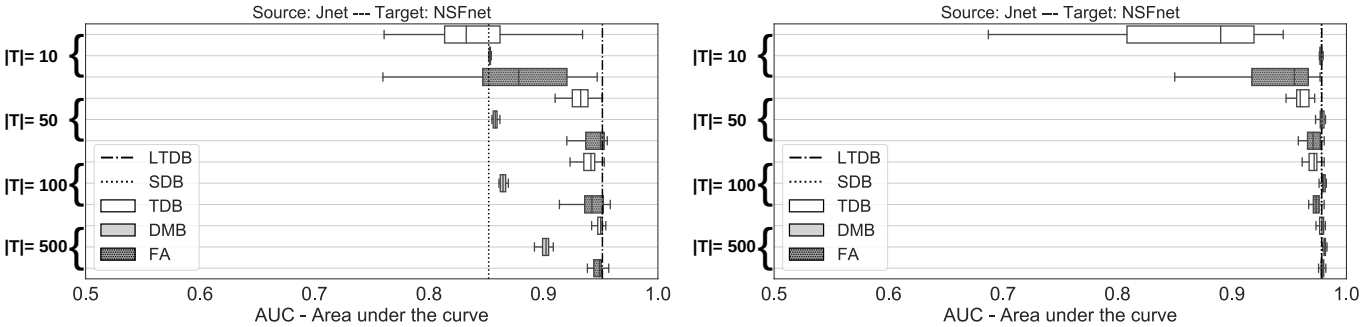

Fig. 11. AUC comparison using Jnet ($\alpha = 1$) as source domain and NSFnet ($\alpha = 0.25$) as target domain dataset, using Logistic Regression (left) and Random Forests (right).

techniques in optical networks," IEEE Commun. Surv. & Tutorials **21**, 1383–1408 (2018).

2. K. Christodoulopoulos, P. Kokkinos, A. Di Giglio, A. Pagano, N. Argyris, C. Spatharakis, S. Dris, H. Avramopoulos, J. Antona, C. Delezoide, P. Jennevé, J. Pesic, Y. Pointurier, N. Sambo, F. Cugini, P. Castoldi, G. Bernini, G. Carrozzo, and E. Varvarigos, "Orchestra-optical performance monitoring enabling flexible networking," in *2015 17th International Conference on Transparent Optical Networks (ICTON),* (IEEE, 2015), pp. 1–4.
3. L. Velasco, B. Shariati, F. Boitier, P. Layec, and M. Ruiz, "Learning life cycle to speed up autonomic optical transmission and networking adoption," J. Opt. Commun. Netw. **11**, 226–237 (2019).
4. D. Azzimonti, C. Rottondi, and M. Tornatore, "Reducing probes for quality of transmission estimation in optical networks with active learning," J. Opt. Commun. Netw. **12**, A38–A48 (2020).
5. J. Yu, W. Mo, Y.-K. Huang, E. Ip, and D. C. Kilper, "Model transfer of QoT prediction in optical networks based on artificial neural networks," IEEE/OSA J. Opt. Commun. Netw. **11**, C48–C57 (2019).
6. L. Xia, J. Zhang, S. Hu, M. Zhu, Y. Song, and K. Qiu, "Transfer learning assisted deep neural network for osnr estimation," Opt. Express **27**, 19398–19406 (2019).
7. Q. Yao, H. Yang, A. Yu, and J. Zhang, "Transductive transfer learning-based spectrum optimization for resource reservation in seven-core elastic optical networks," J. Light. Technol. **37**, 4164–4172 (2019).
8. S. Sun, H. Shi, and Y. Wu, "A survey of multi-source domain adaptation," Inf. Fusion **24**, 84 – 92 (2015).
9. R. Di Marino, C. Rottondi, A. Giusti, and A. Bianco, "Assessment of domain adaptation approaches for QoT estimation in optical networks," in *2020 Optical Fiber Communications Conference and Exhibition (OFC),* (Optical Society of America, 2020), p. Th3D.2.
10. Y. Cheng, W. Zhang, S. Fu, M. Tang, and D. Liu, "Transfer learning simplified multi-task deep neural network for pdm-64qam optical performance monitoring," Opt. Express **28**, 7607–7617 (2020).
11. X. Chen, B. Li, R. Proietti, C.-Y. Liu, Z. Zhu, and S. B. Yoo, "Demonstration of distributed collaborative learning with end-to-end qot estimation in multi-domain elastic optical networks," Opt. Express **27**, 35700–35709 (2019).
12. C.-Y. Liu, X. Chen, R. Proietti, and S. B. Yoo, "Evol-tl: evolutionary transfer learning for QoT estimation in multi-domain networks," in *2020 Optical Fiber Communications Conference and Exhibition (OFC),* (Optical Society of America, 2020), p. Th3D.1.
13. Z. Xu, C. Sun, T. Ji, H. Ji, and W. Shieh, "Transfer learning aided neural networks for nonlinear equalization in short-reach direct detection systems," in *2020 Optical Fiber Communications Conference and Exhibition (OFC),* (Optical Society of America, 2020), p. T4D.4.
14. X. Chen, R. Proietti, C.-Y. Liu, Z. Zhu, and S. B. Yoo, "Exploiting multi-task learning to achieve effective transfer deep reinforcement learning in elastic optical networks," in *Optical Fiber Communication Conference and Exhibition (OFC),* (Optical Society of America, 2020), p. M1B.3.
15. W. L. Hamilton, R. Ying, and J. Leskovec, "Representation learning on graphs: Methods and applications," arXiv preprint arXiv:1709.05584 (2017).
16. N. M. Kriege, F. D. Johansson, and C. Morris, "A survey on graph kernels," Appl. Netw. Sci. **5**, 1–42 (2020).
17. H. Daumé III, "Frustratingly easy domain adaptation," arXiv preprint arXiv:0907.1815 (2009).
18. B. Sun, J. Feng, and K. Saenko, "Correlation alignment for unsupervised domain adaptation," in *Domain Adaptation in Computer Vision Applications,* (Springer, 2017), pp. 153–171.
19. C. Rottondi, L. Barletta, A. Giusti, and M. Tornatore, "Machine-learning method for quality of transmission prediction of unestablished lightpaths," J. Opt. Commun. Netw. **10**, A286–A297 (2018).