An Open EDFA Gain Spectrum Dataset and Its Applications in Data-driven EDFA Gain Modeling

Zehao Wang[1,*], Dan Kilper[2], and Tingjun Chen[1]

[1]*Department of Electrical and Computer Engineering, Duke University, Durham, North Carolina 27708, USA*
[2]*CONNECT Centre, Trinity College Dublin, Dublin, Ireland*
[*]*Corresponding author: zehao.w@duke.edu*

Compiled April 24, 2023

Optical networks satisfy high bandwidth and low latency requirements for telecommunication networks and data center interconnection. To improve the network resource utilization, machine learning (ML) is used to accurately model optical amplifiers such as erbium-doped fiber amplifiers (EDFAs) which impact end-to-end system performance such as quality of transmission (QoT). However, a comprehensive measurement dataset is required for ML to accurately predict an EDFA's wavelength-dependent gain. We present an open dataset consisting of 202,752 gain spectrum measurements collected from 16 commercial-grade ROADM booster and pre-amplifier EDFAs under varying gain settings and diverse channel loading configurations over 2,785 hours in total, with a total dataset size of 3.1 GB. With this EDFA dataset, we implemented component-level deep neural network (DNN) based EDFA models and use transfer learning (TL) to transfer the EDFA model among 16 ROADM EDFAs, which achieve less than 0.18/0.24 dB mean absolute error for booster/pre-amplifier gain prediction using only 0.5% of the full target training set. We also showed that TL reduces the EDFA data collection requirements on a new gain setting or a different type of EDFA on the same ROADM.

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

1. INTRODUCTION

Telecommunication networks and cloud infrastructure rely on amplified optical networks to deliver high data rates over metro and long-haul distances. Reconfigurable optical add-drop multiplexers (ROADMs) are used to add and drop signals within such networks, and erbium-doped fiber amplifiers (EDFAs), sometimes in tandem with Raman amplifiers, are used to overcome node and link losses. The EDFA output power is typically the main determinant of the signal launch power and the EDFA noise figure sets the accumulated amplifier noise levels, which impact end-to-end system performance metrics such as the optical signal-to-noise ratio (OSNR) and other quality of transmission (QoT) measures [1]. However, characterizing the gain spectrum of an EDFA is challenging as it depends on many factors such as the internal hardware architecture, gain setting, channel loading configuration, and input power levels. For these reasons, vendors are motivated to treat the wavelength dependent gain of EDFAs as a variable quantity accounted for through margin allocations. Therefore, better characterization of amplifier gain is of interest to achieving low-margin systems.

Recent work has focused on developing accurate models for the wavelength dependent gain profiles of optical amplifiers like Raman amplifiers [2, 3] and EDFAs [4, 5], which can be further used for effective prediction of the optical power spectrum evolution [6] and QoT estimation [7, 8]. It has been shown that machine learning (ML) models, such as those based on deep neural networks (DNNs) can achieve prediction accuracy primarily limited by the measurement resolution if the model is trained on large EDFA gain spectrum measurement datasets. However, such prior work is built on datasets collected from very few EDFAs and usually only considers a limited set of channel loading configurations and/or input power levels. Moreover, these datasets are not publicly available, therefore making it challenging to compare different EDFA models using the same baseline, as the measurement resolutions and methods may differ from one experiment to the next.

Although the DNN-based EDFA gain model can achieve high gain spectrum prediction accuracy, it requires collecting a comprehensive set of EDFA gain spectrum measurements for *each* EDFA. For example, collecting such a dataset for a single EDFA covering different gain settings and diverse channel

loading configurations can consume up to 51 hours [9]. A promising solution to overcome this challenge is to apply *transfer learning (TL)* [10, 11], which is an ML technique that allows for building a new target model based on a pre-trained source model that shares similar model knowledge using very few data samples collected from the target domain.

In this paper, we make two key contributions aiming to address these challenges. First, we present an open dataset of the gain spectrum measurements for 16 EDFAs within 8 commercial-grade Lumentum ROADM-20 units deployed in the PAWR COSMOS testbed [12]. The dataset includes measurements collected from 8 booster EDFAs, each with 3 gain settings, and 8 pre-amplifier EDFAs, each with 5 gain settings. For each EDFA at a given gain setting, 3,168 gain spectrum measurements are collected with a set of diverse channel loading configurations and varying input power levels. Importantly, all data is collected using the built-in photodiodes (PDs) and optical channel monitors (OCMs) of each ROADM unit, without relying on any external measurement equipment, therefore providing the equivalent of in-situ characterization results. The full 202,752 EDFA gain spectrum measurement dataset collected from 16 EDFAs over 2,785 hours is shared publicly and available at [9]. This dataset can be potentially integrated with emulators such as mininet-optical [13] and planning tools such as GNPy [1]. The dataset serves as an open resource for researchers to evaluate and compare different ML-based EDFA models.

Second, we investigate the use of TL-based EDFA gain models and show that using only 0.5% of the new data collected from the target EDFA (13 measurements), the transferred target model can achieve similar gain prediction accuracy compared to the source model with the full training set (2,732 measurements). We demonstrate three different scenarios that can benefit from TL with a largely reduced EDFA data collection process: (*i*) TL between EDFAs of the same type; (*ii*) TL between different EDFA gain settings, and (*iii*) TL between different EDFA types. For TL between EDFAs of the same type, we achieve an average median absolute error of 0.08 dB for booster amplifiers and 0.10 dB for pre-amplifiers. For TL between different EDFA gain settings, 0.16 dB MAE is achieved averaged from two gain settings transferred to and tested on another gain setting. For TL between different EDFA types, 0.16 dB MAE is achieved. Based on these evaluation results, TL-based EDFA gain models can reduce data measurement times by 200× without sacrificing model prediction accuracy.

The rest of the paper is organized as follows. We review related work in Section 2. We present the EDFA gain spectrum measurement setup and analysis of the collected dataset in Sections 3 and 4. Using the collected dataset, we present the DNN- and TL-based EDFA gain spectrum model in Sections 5 and 6, and conclude in Section 7.

2. RELATED WORK

Traditional EDFA Models. EDFAs in terrestrial wavelength division multiplexed (WDM) systems with channel add-drop multiplexing use automatic gain control (AGC) to maintain a target gain, which controls the total power gain rather than the gain of individual channels. For example, if the target gain setting is 18 dB, the actual gain spectrum can have a fluctuation of ±0.5 dB across different wavelength channels. Moreover, the channel gain deviates from the target gain under different channel loading, input power level, and gain settings, which can be characterized by a physical model given by Eq. (1) [14]

$$\hat{g}(\lambda_i) = \frac{G_{TC}}{G_M} \cdot \left[\frac{\sum_j P_j + N_I + N_C}{\sum_j P_j g_j t_j + g_R N_R + g_I N_I} \right] \cdot g_m(\lambda_i), \quad \textbf{(1)}$$

where G_{TC} and G_M are the target gain and mean gain, respectively. $g_m(\lambda_i)$ is the original channel gain in the i^{th} wavelength channel at λ_i, before new input power P_j and the corresponding residual ripple g_j and the tilt t_j is applied to the j^{th} wavelength channel at λ_j. The noise includes five different factors: the total input noise N_I, total amplifier input-referred noise N_R, amplifier AGC noise compensation factor N_C, average incident noise gain ripple g_I, and input-referred noise gain ripple g_R. However, fully characterizing such factors is challenging due to many practical reasons.

In practice the gain variations described in Eq. (1) follow a center of mass weighting of the channel powers by their wavelength dependent gain functions. Based on this, another well-known model is the center of mass (CM) model, which uses simple measurements to predict the EDFA gain spectrum and for equal channel powers is given by

$$g_{\text{CM}}(\lambda_i) = g_{\text{wdm}}(\lambda_i) + \frac{1}{n} \sum_{j=1}^{n} \left[g_{\text{single}}(\lambda_j) - g_{\text{wdm}}(\lambda_j) \right], \quad \textbf{(2)}$$

where $g_{\text{wdm}}(\lambda_i)$ and $g_{\text{single}}(\lambda_j)$ are the gain of the i^{th} wavelength channel under WDM and single channel loading configurations, respectively. Eq. (2) is usually accurate for the two extreme cases of one channel and all channels turned on, it approximates the gain spectral behavior for other loadings, which can vary significantly for complex multi-stage amplifiers and due to effects such as spectral hole burning [15].

ML-based EDFA Models. Recent research has also focused on using ML to better characterize the wavelength dependent gain spectrum of EDFAs. In particular, a DNN-based EDFA gain model was proposed in [4], where individual sub-models are used to predict the EDFA output power for random channel configurations under one single gain and one tilt setting. The measurement is collected using built-in OCMs and PDs, the first example of an in-situ monitor-based model. Another DNN-based EDFA model with high accuracy to predict partial-fill EDFA gain profile was proposed in [5], which was trained using a dataset consisting of 50,000 measurements using an optical spectrum analyzer (OSA) with high resolution. The prediction was related to the WDM measurements as the output of the model was the power difference between fully loaded (WDM) and partially loaded (arbitrary) channel power. [16] considers optical signal-to-noise ratio (OSNR) prediction using EDFA models that use two different models to predict gain profile and noise figure, separately, with an additional OSA for data collection.

Although individual EDFA models were well investigated, there are still limitations regarding the use of the ML-based approach. First, training ML-based EDFA models requires tremendous training data that is time-consuming to collect in real-world scenarios, especially in already deployed networks. Second, an EDFA model only applies to itself, and data recollection and model retraining are required for new EDFAs.

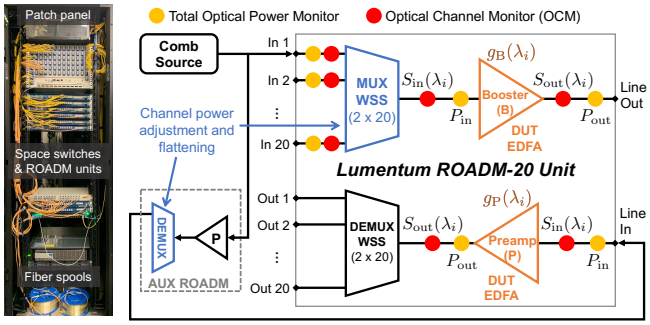

Fig. 1. (Left) COSMOS optical data center with ROADM devices. (Right) block diagram of the Lumentum ROADM-20 unit and the measurement setup for the DUT booster/pre-amplifier EDFA.

[17] proposed a hybrid ML-based EDFA model that combines an analytical model and an ML-based model to reduce the training dataset size and training time. The results showed that 46% training samples or 20% training time were reduced when using the output from an analytical model feeding into the ML model. This method reduced data collection time but still train each EDFA model individually. In addition, [18] showed that a pre-trained ML-based EDFA model can be directly extended to multiple physical devices of the same make with small prediction error by training an EDFA model using measurements collected from multiple EDFAs with a benchtop OSA. However, this approach assumed EDFAs' gain profiles of the same make were highly similar as it used a single model to predict multiple EDFA gain profiles.

ML-based EDFA models were also integrated with multi-span optical transmission systems for QoT prediction. In [19], a generalized EDFA model trained on separately collected gain spectrum measurements using an OSA is used to predict the OSNR across 8 channels in a 3-span link. In [20], OSNR prediction in a 20-span link with 40 channels using characterized inline EDFAs is demonstrated, without considering the model generalization to different topologies. In [6], individually characterized component-level EDFA models were applied to 5-span ROADM systems with 95 channels and 10 EDFAs, where each model is trained using measurements collected using built-in OCMs and photodiodes of the ROADM units.

Compared to this prior art, our work focuses on creating an EDFA open dataset and component-level EDFA modeling. We collected gain profile measurements on 16 commercial-grade Lumentum EDFA devices with different channel loading configurations, input power levels, and gain settings. We report the gain profile difference for EDFAs of the same make and show gain profile variations across a long time period. In addition to the EDFA dataset, we also show measurement time can be largely reduced with TL for EDFAs of the same make. In particular, TL can be used between different EDFA devices of the same type, different gain settings on the same EDFA device, and different EDFA types on the same ROADM. A portion of this paper is an expansion of our recent work [21].

3. EDFA GAIN SPECTRUM MEASUREMENT SETUP AND DATA COLLECTION

We now describe the EDFA gain spectrum measurement setup using the COSMOS testbed and the data collection pipeline.

A. The PAWR COSMOS Testbed

The PAWR COSMOS testbed is a city-scale optical-wireless programmable testbed being deployed in Manhattan, New York City, to support advanced optical and wireless experiments [22]. A more detailed description about COSMOS' programmable optical network and the supported applications can be found in [12]. In particular, the testbed consists of one Calient S320 320×320 space switch, one Dicon 16×16 space switch, 8 commercial-grade Lumentum ROADM units, one customized comb source, various lengths of fiber spools, and a dark fiber network between Columbia University, the colocation facility at 32 Avenue of the Americas (32 AoA), and the City College of New York (CCNY), some of which is shown in Fig. 1. Using the space switching and WDM switching capabilities, different topologies in the optical physical layer that emulate varying metro networks can be constructed [6, 23].

B. EDFA Gain Spectrum Measurement Setup

We characterize the gain spectrum of 16 EDFAs of two types: *booster (B)* and *pre-amplifier (P)*, as part of 8 commercial-grade Lumentum ROADM-20 units. Fig. 1 shows the block diagram of the Lumentum ROADM-20 unit and the measurement setup of a device under test (DUT) EDFA. Each ROADM unit consists of one MUX wavelength selective switch (WSS), one DEMUX WSS, one booster EDFA (at line out), and one pre-amplifier EDFA (at line in). Each ROADM is also equipped with total power and channel power monitoring capabilities using the built-in PDs and OCMs with a power measurement resolution of 0.01 dB and 0.1 dB, respectively. We use a comb source to generate a set of 95×50 GHz WDM channels in the C-band following the ITU DWDM 50 GHz grid specification [24].

Fig. 1 shows the booster and pre-amplifier EDFA measurement topology. With a DUT booster EDFA, the output of the comb source is connected to an add port of the MUX WSS, which applies the channel loading configuration, adjusts the power level in each loaded channel, and generates a flat input power spectrum to the DUT EDFA. The output of the DUT booster is terminated. Similarly, with a DUT pre-amplifier EDFA, the output of the comb source is first connected to the pre-amplifier EDFA and DEMUX WSS of an auxiliary ROADM, whose DEMUX WSS applies the channel loading configuration, adjusts the power level in each loaded channel, and generates a flat output power spectrum at the input of the DUT pre-amplifier EDFA. The output of the DUT pre-amplifier EDFA is terminated by the following DEMUX WSS. The wavelength dependent gain spectrum of each EDFA, denoted by $g(\lambda_i)$, can be characterized by its input power spectrum, $S_{\text{in}}(\lambda_i)$, and output power spectrum, $S_{\text{out}}(\lambda_i)$, i.e.,

$$g(\lambda_i) = S_{\text{out}}(\lambda_i) - S_{\text{in}}(\lambda_i), \forall i = 1, 2, \ldots, 95, \quad \textbf{(3)}$$

with $\lambda_1 = 1,529.16\,\text{nm}$ (196.050 THz) and $\lambda_{95} = 1,566.72\,\text{nm}$ (191.350 THz).

We use the Network Configuration Protocol (NETCONF) with Yet Another Next Generation (YANG) data modeling language to control and collect data from each Lumentum ROADM unit. For example, we use the `add-connection` command to apply the channel loading configuration, whose input parameters include the MUX/DEMUX WSS module, connection index, start/end frequencies, attenuation, input/output ports, and channel block status. After waiting for a certain

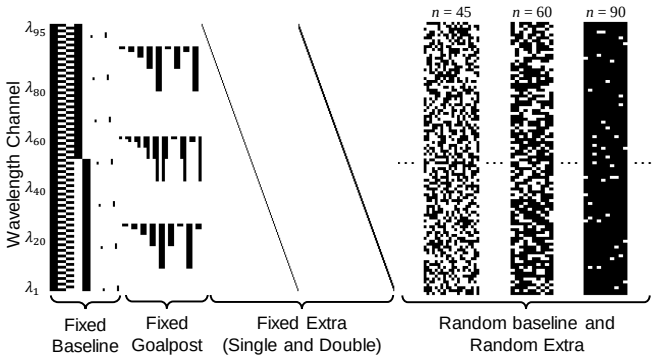

Fig. 2. The collected EDFA gain spectrum measurements based on diverse channel loading configurations.

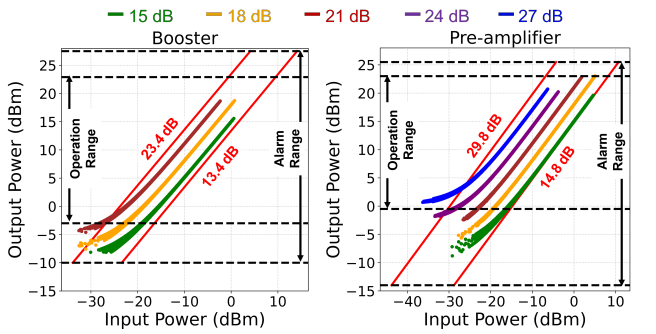

Fig. 3. Input-output power of the collected EDFA gain spectrum measurements overlaid on the EDFA gain masks (high gain mode).

amount of time for the optical system to stabilize, we use the `monitored-channels` and `monitored-connections` commands to obtain the input and output channel power spectrum measurements. EDFA-related information, such as the gain setting and gain tilt, can be retrieved via the `edfas` command. The collected EDFA input/output power spectrum together with other system information is stored in machine-actionable `json` files, which we describe in Section 3-D.

C. Channel Loading Configurations

One main challenge associated with the data collection process is the large number of channel loading configurations, which can affect the wavelength dependent gain, $g(\lambda_i)$, of each EDFA. However, it is impossible to measure all 2^{95} configurations with 95×50 GHz channels where each channel can be switched ON/OFF and all different input channel power levels. To address these challenges, we carefully design 5 sets of diverse channel loading configurations (see Fig. 2) with different numbers of channels n:

(i) **Fixed Baseline** includes the fully loaded (WDM) (with $n = 95$), 4 half loaded (lower/upper/even/odd with $n \in \{47,48\}$), and 7 selected single/double (adjacent) loaded (with $n \in \{1,2\}$) channel configurations;

(ii) **Fixed Goalpost** focuses on two sets of consecutive channels located in 3 channel groups (with short/medium/-long wavelength), and includes 15 balanced and 12 imbalanced goalpost channel configurations with $n \in \{2,4,8,16,32\}$ and $n \in \{9,18\}$, respectively;

(iii) **Fixed Extra** includes the complete set of 95 single and 94 double (adjacent) channel loading configurations;

(iv) **Random Baseline** includes 100, 50, 20 random channel loading configurations for each value of $n \in \{1,2,\ldots,5\}$, $\{6,8,\ldots,20\}$, $\{21,24,\ldots,48\}$, respectively;

(v) **Random Extra** expands *Random Baseline* and includes 10 random channel loading configurations for each value of $n \in \{1,2,\ldots,94,95\}$.

D. Collected Dataset

We consider a target gain of $g_{\mathrm{B}} \in \{15,18,21\}$ dB and $g_{\mathrm{P}} \in \{15,18,21,24,27\}$ dB for each booster and pre-amplifier EDFA, respectively, in the high gain mode with 0 dB gain tilt. For each of the 16 EDFAs at a given gain setting, a total number of 3,168 measurements are collected, where for each channel loading configuration, we also collected repeated measurements with varying EDFA input power levels for comprehensiveness, as summarized in Table 1. In particular, each measurement is stored in the machine-actionable `json` COSMOS EDFA format (see Listing 1 for the structure of the captured measurement data), which includes:

(*i*) The input and output power spectrum of the EDFA measured by the OCM, $S_{\mathrm{in}}(\lambda_i)$ and $S_{\mathrm{out}}(\lambda_i)$, from which $g(\lambda_i)$ can be derived;

(*ii*) The total input and output power of the EDFA measured by the PD, P_{in} and P_{out};

(*iii*) Auxiliary information such as the EDFA gain setting, channel loading configuration, and WSS attenuation setting.

The EDFA gain profile measurement can be time-consuming, mainly due to the time it takes to set the WSS attenuation values (0.85 seconds) and channel loading configuration (3 seconds), and to fetch the OCM/PD readings (6 seconds). To guarantee that the OCM power readings are reliable, extra waiting time is applied depending on channel loading conditions. On average, each measurement lasts ∼41 seconds and ∼58 seconds for the booster and pre-amplifier EDFA, respectively. **Overall, the collected dataset with a size of 3.1 GB includes a total number of 202,752 gain spectrum measurements across 16 EDFAs collected over 2,785 hours**.

4. EDFA GAIN SPECTRUM MEASUREMENT RESULTS

We now provide a quantitative overview of the collected EDFA gain profile measurement dataset. For each EDFA with a given channel loading configuration, we focus on:

(i) The *total input-output power relationship*, $(P_{\mathrm{in}}, P_{\mathrm{out}})$;

(ii) The *gain ripple* as a function of the wavelength, given by $\widetilde{g}(\lambda_i) = g(\lambda_i) - g_0, \forall i$, where g_0 is the EDFA gain setting, i.e., we consider gain ripple *normalized to the target gain* instead of with zero mean;

(iii) The *peak-to-peak gain ripple*, given by $\max_i\{\widetilde{g}(\lambda_i)\} - \min_i\{\widetilde{g}(\lambda_i)\}$ across the loaded wavelength channels.

Fig. 3 shows the total input-output power relationship, $(P_{\mathrm{in}}, P_{\mathrm{out}})$, of the collected EDFA gain spectrum measurement dataset obtained using the built-in PDs (see Fig. 1). The input-output power measurements are also overlaid on the EDFA gain mask (in the high-gain mode) with the corresponding operation/alarm range of each EDFA specified by the vendor (Lumentum). In particular, different curves represent the measured $(P_{\mathrm{in}}, P_{\mathrm{out}})$ values under different gain settings across all EDFAs of the same type. It can be seen that the collected dataset covers a significant portion of the high gain range for both the booster EDFA (13.4–23.4 dB) and pre-amplifier EDFA (14.8–29.8 dB). Overall, most of the measurement results exhibit a linear input-output power relationship, except for scenarios where the EDFA output power is close to the

Table 1. Summary of the measurements for each EDFA (booster or pre-amplifier) under a given gain setting.

Channel loading configurations	(# of loaded channels, n) × (# of channel loading config. with the same n)	# of different input channel power levels	# of repeated meas. for each channel loading config.	Total # of meas.
Fixed Baseline	95×1 (WDM), 48×2 (lower/odd), 47×2 (upper/even), 1×7 (single), 2×7 (double)	5 (WDM), 2 (single/double), 2 (lower/upper/odd/even)	20 (WDM), 5 (single/double), 5 (lower/upper/odd/even)	280
Fixed Goalpost	$[2, 4, 8, 16, 32] \times 3$ (balanced), $[9, 18] \times 6$ (imbalanced)	2 (balanced), 2 (imbalanced)	5 (balanced), 5 (imbalanced)	270
Fixed Extra	1×95 (single), 2×94 (double)	2 (single), 2 (double)	2 (single), 1 (double)	568
Random Baseline	$[1, 2, \ldots, 5] \times 100$, $[6, 8, \ldots, 20] \times 50$, $[21, 24, \ldots, 48] \times 20$	1	1	1,100
Random Extra	$[1, 2, \ldots, 94, 95] \times 10$	1	1	950

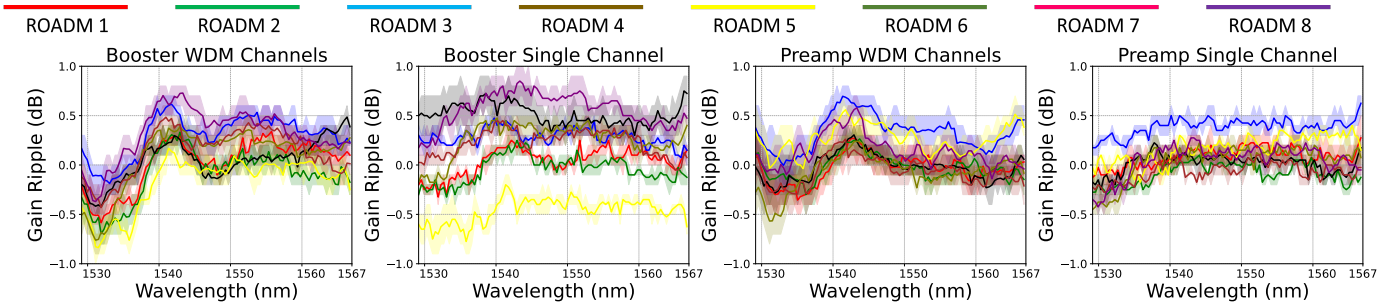

Fig. 4. Measured gain ripple of 16 EDFAs at 18 dB gain setting, under WDM/single channel loading configurations.

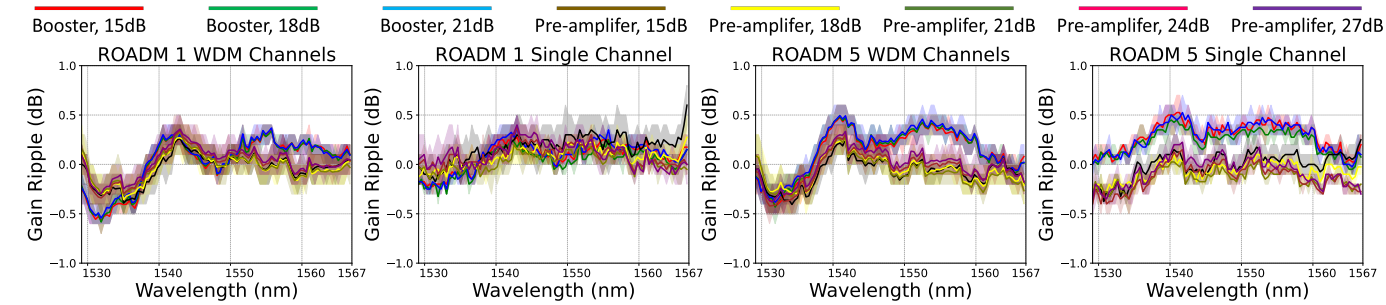

Fig. 5. Measured gain ripple of 4 EDFAs at different gain settings, under WDM/single channel loading configurations.

lower limit of the operation range, e.g., $P_{\text{in}} < -20$ dBm. These measurements are important since the built-in OCMs can be ensured to maintain the 0.1 dB measurement resolution when the EDFA operates within the operation range, and may provide alarming and fault detection when the EDFA operates outside the operation range but inside the alarm range.

Through analysis of the gain ripple spectrum of individual EDFAs under different gain settings and channel loading configurations, a better understanding of the wavelength dependent gain spectrum among all tested EDFAs can be derived. Fig. 4 shows examples of the measured gain ripple spectrum, $\tilde{g}(\lambda_i)$, for all 16 EDFAs at 18 dB gain setting, under the full (WDM) and single channel loading configurations. The gain ripple is normalized to the target gain (instead of with zero mean) and the different gain profiles across EDFAs can be clearly visualized. The solid lines represent the mean gain ripple averaged across all measurements, while the shaded areas represent the full range of the measured gain ripple, including the minimum and maximum values. It can be seen that different types of EDFA (booster or pre-amplifier), or different EDFA devices of the same type (e.g., booster EDFA of two ROADM units), have different gain ripple spectra and that the gain ripple spectrum of each EDFA also depends on the channel loading configurations (WDM or single channels).

In addition, Fig. 5 shows the measured gain ripple spectra of the EDFAs in ROADM 1 and ROADM 5, across all considered gain settings and under the full (WDM) and single channel loading configurations. It can be seen that with the full (WDM) channel loading configuration, the gain ripple spectrum for the booster EDFA is similar across the 15/18/21 dB gain settings, but are different from that of the pre-amplifier EDFA across the same gain settings, especially in channels with longer wavelengths. Similarly, the gain ripple spectrum for the pre-amplifier EDFA is similar across the 15/18/21/24/27 dB gain settings. Overall, it can be observed that the gain ripple profile for each EDFA depends on many factors including the gain setting, channel loading configurations, and input power level. To obtain overall statistics of the gain ripple spectrum for different EDFA devices, types, and gain settings, Fig. 6 shows the mean peak-to-peak gain ripple averaged across all channel loading configurations for each EDFA at a given gain setting. In particular, each entry represents the mean peak-to-peak gain ripple for one EDFA at a given gain setting, averaged across 3,168 gain spectrum measurements. The results show that the mean peak-to-peak gain ripple is within a range of 0.5–0.9 dB but varies by different EDFA types, devices, and gain settings.

We also evaluate the variation of the EDFA gain spectrum

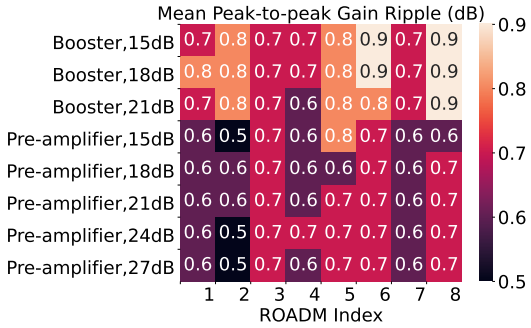

Fig. 6. Mean peak-to-peak gain ripple across 8 booster EDFAs and 8 pre-amplifier EDFAs with different gain settings.

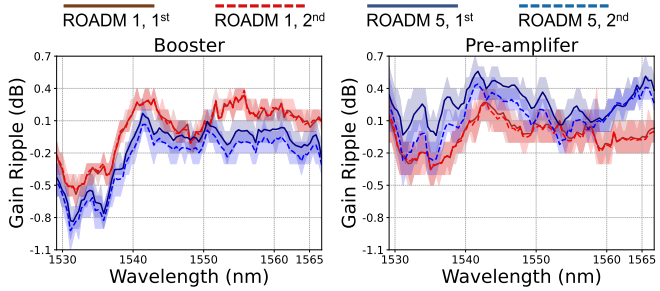

Fig. 7. Example gain ripple spectrum measurements of the booster and pre-amplifier EDFAs at 18 dB gain setting and under WDM channel loading configurations, spanning 10 months.

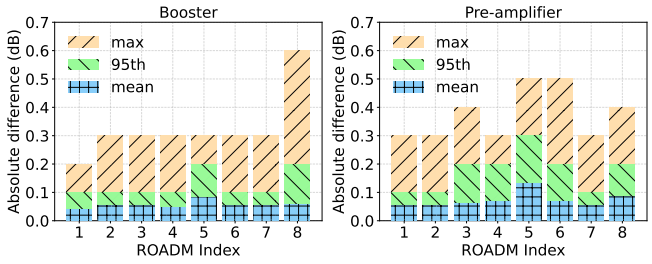

Fig. 8. Mean, 95$^{\text{th}}$-percentile, and maximum values of the absolute difference in the two rounds of EDFA gain ripple spectrum measurements spanning 10 months.

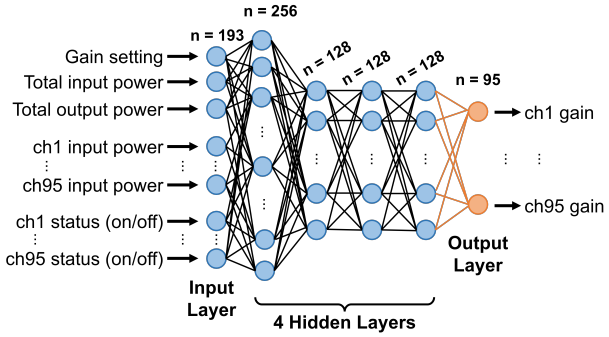

Fig. 9. Architecture of the deep neural network (DNN) model used for EDFA gain prediction. Transfer learning reinitializes the orange output layer before retraining.

across a long time period, which is important due to factors such as the potential aging of the hardware. In particular, after 10 months since the completion of the initial dataset collection, we re-collect the gain spectrum measurements for all 16 EDFAs at 18 dB gain settings and under the same channel loading configurations (see Table 1). Fig. 7 shows the gain ripple spectrum for the EDFAs in ROADMs 1 and 5 under the WDM (full) channel loading configurations from the first and second measurement round using solid and dashed lines, respectively. It can be seen that the difference in the gain ripple spectrum across the measurements spanning 10 months is minimum, i.e., the difference in the mean gain ripple is only $<$0.2 dB. We also analyze the mean/95$^{\text{th}}$-percentile/maximum absolute difference in the gain spectrum measurements spanning 10 months for individual booster and pre-amplifier EDFAs, and the results are shown in Fig. 8. Note that the 95$^{\text{th}}$-percentile and maximum absolute difference has a resolution of 0.1 dB, due to the 0.1 dB measurement resolution of the built-in OCMs. Overall, the mean difference in the gain spectrum measurements spanning 10 months is less than 0.1 dB, while the 95$^{\text{th}}$-percentile difference is within 0.3 dB. We would like to note that more measurements are ongoing, with the aim to provide a comprehensive characterization of the EDFA gain spectrum across different devices and time spans.

5. DNN-BASED EDFA GAIN SPECTRUM MODEL

In this section, we present a DNN-based EDFA model for characterizing the wavelength dependent gain spectrum using the collected dataset, and compare it against the CM model.

A. Architecture of the DNN-based EDFA Model

Fig. 9 shows the DNN model architecture, which consists of an input layer, four hidden layers with 256/128/128/128 neurons, and an output layer, where the neurons are initialized by the Kaiming normalization. The input features to the DNN model include the EDFA gain setting (g_0), total input and output power (P_{in} and P_{out}), input power spectrum ($S_{\text{in}}(\lambda_i)$), and a binary vector indicating the channel loading configuration, denoted by $\mathbf{c} = [c_i]_{i=1}^{95}$ with

$$c_i = \begin{cases} 1, & \text{if the } i^{\text{th}} \text{ wavelength channel is switched on,} \\ 0, & \text{otherwise.} \end{cases}$$

The output layer predicts the EDFA gain spectrum, $g(\lambda_i)$, corresponding to the input parameters. For the input and hidden layers, we apply batch normalization and use the exponential linear unit (ELU) activation function. We consider the following loss function based on the mean squared error (MSE) of the predicted and ground truth gain spectrum profile across the loaded channels, i.e.

$$MSE = \frac{1}{\sum_{i=1}^{95} c_i} \cdot \sum_{i:c_i=1} \left[g_{\text{pred}}(\lambda_i) - g_{\text{meas}}(\lambda_i)\right]^2, \qquad \textbf{(4)}$$

where $g_{\text{pred}}(\lambda_i)$ and $g_{\text{meas}}(\lambda_i)$ denote the predicted and measured gain in the i^{th} wavelength channel, respectively. A component-level DNN model is trained for each EDFA with the same setting: a gradient clipping threshold of 3.0 and a learning rate of 0.001 over 600 epochs.

We use all the EDFA gain spectrum measurements under three gain settings of 15/18/21 dB to train and test the performance of the DNN-based EDFA gain model. Note that although there are two additional gain settings for the pre-amplifier (24/27 dB), we only choose the dataset corresponding to the 15/18/21 dB gain setting to keep a consistent size of the dataset used for the DNN model training and testing across the booster and pre-amplifier EDFA types. For each gain setting, we split the EDFA gain measurement dataset

Table 2. Dataset split for training (including both the DNN- and TL-based EDFA gain models) and test at each EDFA gain setting.

	Training set	Test set (baseline)	Test set (goalpost)	Training set for TL
# of meas.	2,732	220	216	13
Ratio	86%	7%	7%	0.5%

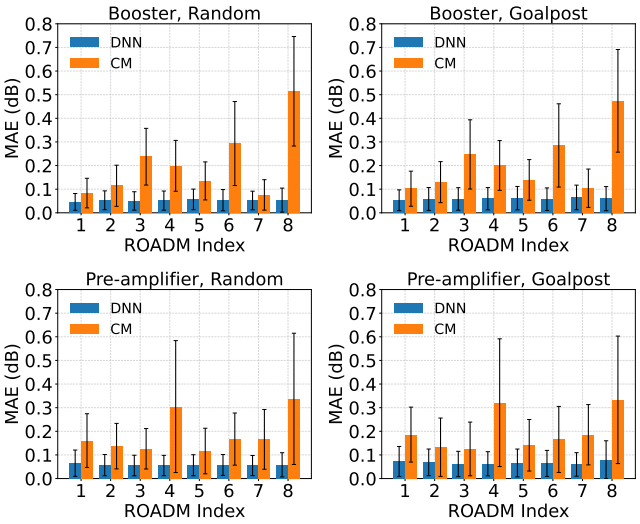

Fig. 10. Mean absolute error (MAE) of component DNN and CM EDFA model across 8 boosters and 8 pre-amplifiers on two test sets.

into the training/test sets with a split ratio of 0.86/0.14: 2,732 gain spectrum measurements are used as the training set, and the remaining 436 gain spectrum measurements are used as the test set. Specifically, the test set includes 20% of the *Fixed Goalpost* (216 measurements) and 20% of the *Random Baseline* (220 measurements) gain spectrum measurements, which represent a diverse set of channel loading configurations with randomly selected channels and groups of close-by channels (see Table 2).

B. Performance of the DNN-based EDFA Gain Model

We now show the performance of the developed DNN-based model and compare it with the CM model (Eq. (2)). Fig. 10 shows the mean absolute error (MAE) and standard deviation of the EDFA gain spectrum predicted by the component-level DNN and CM models, across 8 booster and pre-amplifier EDFAs using test sets with different channel loading configurations (random and goalpost). Specifically, across the 8 booster EDFAs, the DNN model achieves an average MAE of 0.05 dB and 0.06 dB under the random and goalpost test set, respectively, outperforming that achieved by the CM model (0.21 dB and 0.21 dB, respectively). Similarly, across the 8 pre-amplifier EDFAs, the DNN model achieves an average MAE of 0.06 dB and 0.07 dB under the random and goalpost test set, respectively, again outperforming that achieved by the CM model (0.19 dB and 0.20 dB, respectively). The DNN-based model also achieves a smaller variance compared to the CM model. This demonstrates that the DNN model is able to improve the EDFA gain spectrum prediction accuracy via learning from a large measurement dataset. Overall, the DNN-based model suffers from slightly larger prediction errors for the goalpost test set compared to the random test

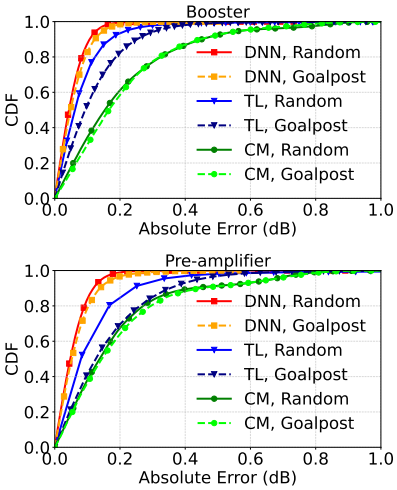

Fig. 11. CDF of absolute errors on component DNN and CM EDFA models across 8 booster EDFAs and 8 pre-amplifier EDFAs.

Table 3. Maximum absolute error for EDFA gain spectrum prediction achieved by the CM and DNN models.

Booster	Random	Goalpost	Pre-amplifier	Random	Goalpost
CM	1.13 dB	1.09 dB	CM	1.14 dB	1.08 dB
DNN	0.81 dB	0.50 dB	DNN	0.89 dB	0.57 dB

set.

Fig. 11 shows the cumulative distribution function (CDF) of the absolution error of gain spectrum prediction across all EDFAs of the same type, under both the DNN and CM model. The results show that for booster EDFAs, the DNN model is able to achieve a median gain prediction error of 0.05/0.06 dB for the random/goalpost test set, which is significantly smaller than that achieved by the CM model (0.21/0.21 dB for the random/goalpost test set). In terms of the tail performance, the DNN model is able to achieve a 95^{th}-percentile gain prediction error of 0.13/0.15 dB for the random/goalpost test set, which is again much smaller than that achieved by the CM model (0.58/0.58 dB for the random/goalpost test set). The results for the pre-amplifiers show similar trends when comparing the performance achieved by the DNN and CM models. In addition, Table 3 shows the maximum absolute errors of the gain spectrum prediction achieved by the DNN and CM models. For both test sets, the max absolute errors of DNN are smaller than those achieved by the CM models.

6. TL-BASED EDFA GAIN SPECTRUM MODEL

Transfer learning (TL) is an ML method that uses domain knowledge from a pre-trained model to apply to a new but similar problem. In this section, we show that TL can be applied for modeling the gain spectrum modeling across different EDFAs with minimum data collection.

A. TL Model and Target Dataset Selection

We apply the following TL procedure to transfer a DNN-based *source model* to a *target model*. First, the input layer and all four hidden layers of the DNN (see Fig. 9) are frozen, which are treated as the feature extractor of the DNN model, and the weights of the output layer using the Kaiming normalization are reinitialized. Then, the DNN model is re-trained using

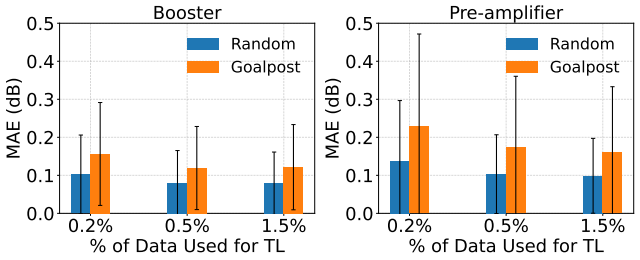

Fig. 12. MAE of the EDFA gain spectrum prediction achieved by the TL-based EDFA gain model with varying target to source data size ratios, $N_{\text{tgt}}/N_{\text{src}}$.

the same MSE loss function given by Eq. (4) with a step size of 0.05 over 150 epochs. Finally, all layers are unfrozen and fine-tuned with a step size of 0.001 over 20 epochs, while the batch normalization parameters are kept unchanged.

Using the dataset and DNN-based EDFA model described above, we first investigate for a given (pre-trained) source model, how much new data is needed from a target EDFA. We consider all cases where each booster/pre-amplifier EDFA serves as the source model, which is then transferred to each of the 7 other booster/pre-amplifier EDFAs using different sizes of target EDFA datasets. Let N_{tgt} and N_{src} denote the number of gain spectrum measurements at each gain setting used to train the source model and to construct the target model for TL, respectively. We consider $N_{\text{src}} = 2,732$ gain spectrum measurements from the source EDFA and different numbers of measurements from the target EDFA:

(i) $N_{\text{tgt}} = 5$ gain spectra under fully loaded channel configurations, with $N_{\text{tgt}}/N_{\text{src}} = 0.2\%$;

(ii) $N_{\text{tgt}} = 13$ gain spectra under fully loaded and half loaded channel configurations, with $N_{\text{tgt}}/N_{\text{src}} = 0.5\%$;

(iii) $N_{\text{tgt}} = 41$ gain spectra unsder fully/half/single/double loaded channel configurations, with $N_{\text{tgt}}/N_{\text{src}} = 1.5\%$.

Fig. 12 shows the MAE and standard deviation of the EDFA gain spectrum prediction accuracy averaged across all possible source-target model pairs for the random and goalpost test sets, with varying ratios of $N_{\text{tgt}}/N_{\text{src}}$. The results show that the average EDFA gain prediction accuracy achieved by the target model with a target-source data size ratio of $N_{\text{tgt}}/N_{\text{src}} = 0.5\%$ outperforms that achieved with $N_{\text{tgt}}/N_{\text{src}} = 0.2\%$, but is comparable to that achieved with $N_{\text{tgt}}/N_{\text{src}} = 1.5\%$. Therefore, we empirically select $N_{\text{tgt}} = 13$ with $N_{\text{tgt}}/N_{\text{src}} = 0.5\%$ in the rest of the evaluations, which largely reduces the target data size by 200× while achieving an MAE of <0.2 dB across all EDFAs. Below, we evaluate the performance of TL-based EDFA models in three scenarios.

B. TL between EDFAs of the Same Type

Fig. 13 shows the MAE matrices (in dB) across 8 EDFAs of the same type (booster or pre-amplifier) under the random and goalpost test sets, with three gain settings (15/18/21 dB) and a target data size of $N_{\text{tgt}} = 13$ ($N_{\text{tgt}}/N_{\text{src}} = 0.5\%$). In each MAE matrix, ($i$) entry $(i,i), i = 1,\ldots,8$, corresponds to the component-level DNN-based EDFA model (i.e., without TL), and (ii) entry $(i,j), j \neq i$, corresponds to the transferred EDFA model with the i^{th} and j^{th} EDFA being the source and target model, respectively. For the i^{th} row in the MAE matrix, each entry (i,i) is always smaller than $(i,j), \forall j \neq i$. This shows that the TL-based model always achieves a slightly larger gain

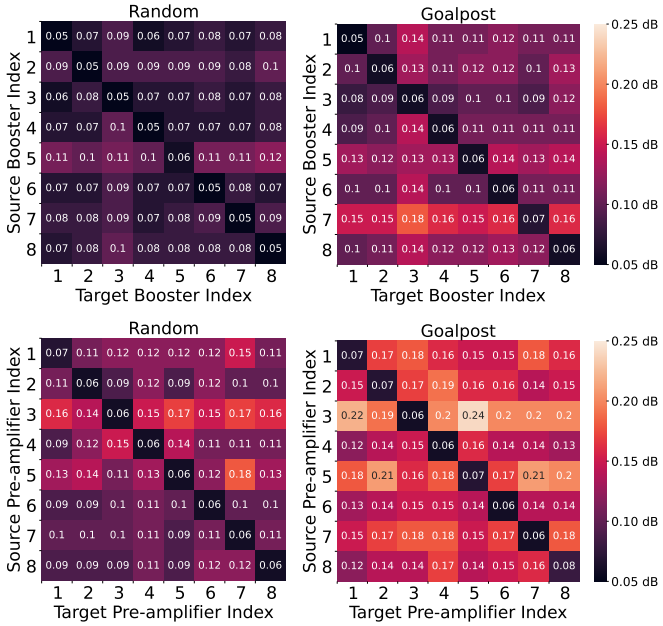

Fig. 13. MAE matrix (in dB) of ML-based EDFA gain spectrum prediction averaged across the random and goalpost test sets, where entry (i,i) corresponds to the DNN-based EDFA model (without TL), and entry $(i,j), i \neq j$ represents TL-based model trained on the i^{th} source EDFA and transferred j^{th} target EDFA.

spectrum prediction error than the DNN-based model without TL, which is as expected given the limited number of new measurements used for deriving the target model.

To compare the TL on two test sets, it can be observed that for booster EDFAs, the TL-based model achieves an MAE between 0.06–0.12 dB and 0.08–0.18 dB on the random and goalpost test set, respectively. Similarly, for pre-amplifier EDFAs, the TL-based model achieves an MAE between 0.09–0.18 dB and 0.12–0.24 dB on the random and goalpost test set, respectively. In particular, TL achieves better average gain prediction accuracy for booster EDFAs compared to the pre-amplifier EDFAs, and suffers from lower accuracy under goalpost channel loading configurations, exhibiting a similar trend as the performance of the component-level DNN model presented in Section 5. We expect that the performance of the target model can be further improved by including (a small number of) gain measurements under the random/goalpost channel loading configurations in the target data. Overall, the MAE achieved by the target booster/pre-amplifier model is within 0.18/0.24 dB across all the test sets.

In addition, Fig. 11 shows the CDF of the absolute prediction error achieved by the TL-based models compared to the DNN-based models. The results show that for booster EDFAs, the TL-based model achieves a median absolute error of 0.06/0.09 dB on random/goalpost test sets, which is (slightly) worse than that achieved by the DNN-based model but outperforms the CM models. Similar trends have been observed for pre-amplifiers EDFAs. In terms of the tail performance, the 95$^{\text{th}}$ absolute errors for booster/pre-amplifier EDFAs achieved by TL-based model prediction are 0.22/0.28 dB and 0.32/0.50 dB for random and goalpost test sets (see Fig. 11). However, TL does not perform well in terms of the maximum absolute error (which can be a few dB) due to the limited new data collected from the target EDFA. Improv-

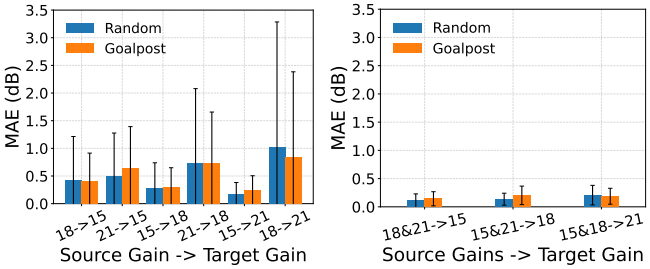

Fig. 14. MAE of the EDFA gain spectrum prediction accuracy using TL from one source gain setting (left) or two source gain settings (right) to another target gain setting on the same booster EDFA.

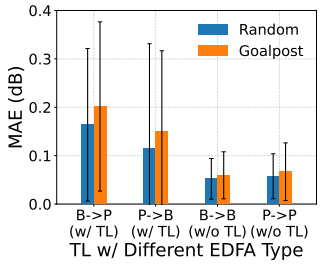

Fig. 15. TL between different EDFA types (B: booster, P: pre-amplifier) on the same ROADM.

ing the prediction accuracy for TL-based EDFA gain models, especially with goalpost channel loading configurations, is considered as a subject of our future research.

C. TL between Gain Settings of the Same EDFA

Fig. 14 shows the MAE and standard deviation of the booster EDFA gain spectrum prediction accuracy achieved by the TL-based model trained using one or two source gain settings, and then transferred to the target model using additional measurements from a new target gain setting. For example, "15 & 18 → 21" means that the source model is trained using EDFA gain spectrum measurements with 15 dB and 18 dB gain settings, and then transferred to the target model using measurements with 21 dB gain setting. The results show that the TL approach using a single source gain setting can result in an MAE of up to 0.8 dB and 1.0 dB under the random and goalpost test set, respectively. These MAE values under the random and goalpost test set can be further reduced to 0.21 dB and 0.19 dB with the additional domain knowledge from the measurements under a second gain setting of 15 dB. In addition, the standard deviation of the gain spectrum prediction error is also reduced with the additional domain knowledge. Overall, the MAE across all source/target gain combinations achieved by the TL-based models with two gain settings is 0.16 dB. Similar MAE performance is observed for TL-based models constructed for the pre-amplifier EDFAs.

D. TL between EDFA Types

So far we consider TL between the *same EDFA type* (booster or pre-amplifier), another way that can benefit the target EDFA data collection process is to apply TL *across different EDFA types*. Fig. 15 shows the MAE and standard deviation averaged across 8 ROADMs of the EDFA gain spectrum prediction accuracy when transferred from a source booster model with three gain settings to a target pre-amplifier model (B→P) or vice versa (P→B) on the same ROADM, compared to the DNN-based model without TL (B→B and P→P). The MAE achieved by the target model is all within 0.21 dB with an average MAE of 0.16 dB, and TL introduces an MAE degradation of only 0.06/0.10 dB and 0.10/0.13 dB for the booster/pre-amplifier EDFA compared to that achieved by the source model under the random and goalpost test set, respectively.

7. CONCLUSIONS

We measured and provided an open EDFA gain spectrum dataset collected from 16 commercial-grade ROADM booster and pre-amplifier EDFAs under different gain settings and channel loading configurations. The dataset includes 202,752 gain spectrum measurements collected over 2,785 hours, with a total dataset size of 3.1 GB. Using this dataset, we investigated TL-based EDFA gain models that can achieve an MAE of less than 0.24 dB using only 0.5% of the full dataset of the new EDFA device. We showed that the EDFA gain models can be transferred between different EDFAs of the same type, different gain settings on the same EDFA, and different EDFA types with varying accuracy.

APPENDIX A: EXAMPLE EDFA GAIN SPECTRUM MEASUREMENT DATASET IN JSON FORMAT

Listing 1. Structure outline for JSON-based dataset files.

```
{"measurement_data": [
    {"open_channel_type": "fully_loaded_channel_wdm",
    "attenuation_setting": 2,
    "calient_input_power_comb_source": 5.83,
    "calient_input_power_roadm_dut_edfa": 6.02,
    "roadm_dut_edfa_info":
        {"input_power": −1.88,
        "output_power": 16.14,
        "target_gain": 18.0,
        "target_gain_tilt": 0.0,
        ...,},
    "roadm_dut_line_port_info": {...},
    "roadm_dut_wss_port_info": {...},
    "roadm_dut_wss_num_active_channel": 95,
    "roadm_dut_wss_active_channel_index": [1,2,...,95],
    "roadm_dut_wss_attenuation":
        {"1": 0.0,
         ...,
         "95": 1.5},
    "roadm_dut_wss_input_power_spectra":
        {"1": −21.8,
         ...,
         "95": −15.9},
    "roadm_dut_wss_output_power_spectra":
        {"1": −26.3,
         ...,
         "95": −21.7},
    "roadm_dut_booster_output":
        {"1": −8.2,
         ...,
         "95": −3.8},},
    ...,]
}
```

Acknowledgements. This work was supported by NSF grants CNS-1827923, OAC-2029295, CNS-2112562, and CNS-2211944, Science Foundation Ireland under Grant #13/RC/2077_P2, a Google Research Scholar Award, and an IBM Academic Award. We thank Ivan Seskar for the help with using the PAWR COSMOS testbed.

REFERENCES

1. A. Ferrari, M. Filer, K. Balasubramanian, Y. Yin, E. Le Rouzic, J. Kundrat, G. Grammel, G. Galimberti, and V. Curri, "GNPy: an open source application for physical layer aware open optical networks," J. Opt. Commun. Netw. **12**, C31–C40 (2020).
2. D. Zibar, A. M. Rosa Brusin, U. C. de Moura, F. Da Ros, V. Curri, and A. Carena, "Inverse system design using machine learning: The raman amplifier case," J. Light. Technol. **38**, 736–753 (2020).
3. A. M. Rosa Brusin, U. C. de Moura, V. Curri, D. Zibar, and A. Carena, "Introducing load aware neural networks for accurate predictions of raman amplifiers," J. Light. Technol. **38**, 6481–6491 (2020).
4. S. Zhu, C. L. Gutterman, W. Mo, Y. Li, G. Zussman, and D. C. Kilper, "Machine learning based prediction of erbium-doped fiber WDM line amplifier gain spectra," in *Proc. ECOC'18,* (2018).
5. Y. You, Z. Jiang, and C. Janz, "Machine learning-based EDFA gain model," in *Proc. ECOC'18,* (2018).
6. Z. Wang, E. Akinrintoyo, D. C. Kilper, and T. Chen, "Optical signal spectrum prediction using machine learning and in-line channel monitors in a multi-span ROADM system," in *Proc. ECOC'22,* (2022).
7. J. Yu, S. Zhu, C. L. Gutterman, G. Zussman, and D. C. Kilper, "Machine-learning-based EDFA gain estimation," J. Opt. Commun. Netw. **13**, B83–B91 (2021).
8. A. Mahajan, K. Christodoulopoulos, R. Martínez, S. Spadaro, and R. Muñoz, "Modeling EDFA gain ripple and filter penalties with machine learning for accurate QoT estimation," J. Light. Technol. **38** (2020).
9. Z. Wang *et al.*, "An open dataset for EDFA gain spectrum measurements," https://wiki.cosmos-lab.org/wiki/Datasets (2022).
10. S. J. Pan and Q. Yang, "A survey on transfer learning," IEEE Transactions on knowledge data engineering **22**, 1345–1359 (2009).
11. L. Torrey and J. Shavlik, "Transfer learning," in *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques,* (IGI global, 2010), pp. 242–264.
12. T. Chen *et al.*, "A software-defined programmable testbed for beyond 5G optical-wireless experimentation at city-scale," IEEE Netw. **36** (2022).
13. B. Lantz, A. A. Díaz-Montiel, J. Yu, C. Rios, M. Ruffini, and D. Kilper, "Demonstration of software-defined packet-optical network emulation with Mininet-Optical and ONOS," in *Proc. IEEE/OSA OFC'20,* (2020).
14. J. Junio, D. C. Kilper, and V. W. Chan, "Channel power excursions from single-step channel provisioning," J. Opt. Commun. Netw. **4**, A1–A7 (2012).
15. K. Ishii, J. Kurumida, and S. Namiki, "Experimental investigation of gain offset behavior of feedforward-controlled wdm agc edfa under various dynamic wavelength allocations," IEEE Photonics J. **8**, 1–13 (2016).
16. Y. You, Z. Jiang, and C. Janz, "OSNR prediction using machine learning-based EDFA models," in *Proc. ECOC'19,* (IET, 2019).
17. S. Zhu, C. Gutterman, A. D. Montiel, J. Yu, M. Ruffini, G. Zussman, and D. Kilper, "Hybrid machine learning EDFA model," in *Proc. IEEE/OSA OFC'20,* (2020).
18. F. Da Ros, U. C. De Moura, and M. P. Yankov, "Machine learning-based EDFA gain model generalizable to multiple physical devices," in *Proc. ECOC'20,* (IEEE, 2020).
19. M. P. Yankov, P. M. Kaminski, H. E. Hansen, and F. Da Ros, "Snr optimization of multi-span fiber optic communication systems employing edfas with non-flat gain and noise figure," J. Light. Technol. **39**, 6824–6832 (2021).
20. S. Kamel, H. Hafermann, D. Le Gac, L. Dos Santos, B. Kégl, Y. Frignac, and G. Charlet, "OSNR prediction for optical links via learned noise figures," in *Proc. ECOC'21,* (2021).
21. Z. Wang, D. Kilper, and T. Chen, "Transfer learning-based ROADM EDFA wavelength dependent gain prediction using minimized data collection," in *Proc. IEEE/OSA OFC'23,* (2023).
22. D. Raychaudhuri, I. Seskar, G. Zussman, T. Korakis, D. Kilper, T. Chen, J. Kolodziejski, M. Sherman, Z. Kostic, X. Gu, H. Krishnaswamy, S. Maheshwari, P. Skrimponis, and C. Gutterman, "Challenge: Cosmos: A city-scale programmable testbed for experimentation with advanced wireless," in *Proc. ACM MobiCom'20,* (2020).
23. E. Akinrintoyo, Z. Wang, B. Lantz, T. Chen, and D. Kilper, "Reconfigurable topology testbeds: A new approach to optical system experiments," Opt. Fiber Technol. **76**, 103243 (2023).
24. T. standardization sector of ITU, "G.694.1 : Spectral grids for WDM applications: DWDM frequency grid," https://www.itu.int/rec/T-REC-G.694.1-202010-I/en (2020).

AUTHOR BIOGRAPHIES

Zehao Wang is currently pursuing his Ph.D. degree in the Department of Electrical and Computer Engineering at Duke University. He received his B.Eng. degree in Information Science from Zhejiang University, China, in 2019, and an M.S. degree in Electrical and Computer Engineering from Carnegie Mellon University in 2020. His research interests are in the areas of next-generation optical-wireless networks with a focus on integrated sensing and communication systems and radio-over-fiber technologies.

Daniel C. Kilper is Professor of Future Communication Networks and SFI CONNECT Centre Director at Trinity College Dublin, Ireland. He holds an adjunct faculty appointment at Columbia University Data Science Institute and College of Optical Sciences, University of Arizona. He is CTO and co-founder of Palo Verde Networks, Inc. and is the Green Internet and Service Provisioning topical area editor for the IEEE Transactions on Green Communications and Networking journal. He co-chairs the IEEE International Network Generations Roadmap (INGR) Optics Working Group. He received MS (1992) and PhD (1996) degrees in Physics from University of Michigan. From 2000-2013, he was a member of technical staff at Bell Labs. His work has been recognized with a NIST Communication Technology Lab Innovator Award, Bell Labs President's Gold Medal Award, and he served on the Bell Labs Presidents Advisory Council on Research. He holds thirteen patents and coauthored six book chapters and more than one hundred seventy peer-reviewed publications. His research is aimed at solving fundamental and real-world problems in communication networks, addressing interdisciplinary challenges for smart cities, sustainability, and digital equity.

Tingjun Chen is an Assistant Professor of Electrical & Computer Engineering and Computer Science at Duke University. He received the B.Eng. degree in electronic engineering from Tsinghua University in 2014, the Ph.D. degree in electrical engineering from Columbia University in 2020, and was a Postdoctoral Associate with Yale University from 2020 to 2021. His research interests are in the area of networking and communications with a specific focus on next-generation wireless networks and Internet-of-Things systems. He received the Google Research Scholars Award, the IBM Academic Award, the Facebook Fellowship, the Wei Family Private Foundation Fellowship, the Columbia Engineering Morton B. Friedman Memorial Prize for Excellence, the Columbia University Eli Jury Award and Armstrong Memorial Award, the ACM SIGMOBILE Doctoral Dissertation Award Runner-Up, and the ACM CoNEXT'16 Best Paper Award.