Knowledge Distillation Applied to Optical Channel Equalization: Solving the Parallelization Problem of Recurrent Connection

Sasipim Srivallapanondh⁽¹⁾, Pedro J. Freire⁽¹⁾, Bernhard Spinnler⁽²⁾, Nelson Costa⁽³⁾, Antonio Napoli⁽²⁾, Sergei K. Turitsyn⁽¹⁾, Jaroslaw E. Prilepsky⁽¹⁾

(1) Aston University, Birmingham, UK; (2) Infinera, Munich, Germany; (3) Infinera, Carnaxide, Portugal s.srivallapanondh@aston.ac.uk

Abstract: To circumvent the non-parallelizability of recurrent neural network-based equalizers, we propose *knowledge distillation* to recast the RNN into a parallelizable feed-forward structure. The latter shows 38% latency decrease, while impacting the Q-factor by only 0.5 dB. © 2022 The Author(s)

1. Introduction

Optical fiber nonlinearity significantly limits the information rate in current coherent transmission systems. Moreover, with the ever-increasing transmission bandwidth, nonlinearity becomes even more important [1]. Various digital signal processing (DSP) techniques have been proposed to minimize non-linear effects [2]. Due to the universal approximation capability of neural networks (NNs), the NNs have recently been intensively studied for the optical channel post-equalization, because they can approximate the inverse optical channel transfer function with good accuracy and revert the nonlinear distortions. In particular, recurrent NNs (RNN) based equalizers have shown the best capability in equalizing nonlinear impairments as compared to the feed-forward NN types [3–5]. However, since the recurrent NN structure has a feedback loop, it is not easily parallelizable. This characteristic makes it challenging to implement the RNNs in low-complexity hardware for high-speed processing [6]. Indeed, parallelization is the key to achieving low latency and, simultaneously, high throughput required by high-speed optical networks, as parallel computing increases the throughput and reduces the computational time [7]. To enable parallelization, in this work, we adopt the concept of knowledge distillation (KD) to transform the RNN-based equalizer into a feedforward structure. Generally, KD is a method to transfer the knowledge from a larger model, known as a *teacher* model, to a more compact one (a *student*), which requires fewer computations [8]. Most of the prior works using KD have focused on classification tasks where the teacher and student have a similar topology. However, KD for a regression task and the cross-architecture KD, as we propose in our work, have only been studied very recently [9]. In this work, we also reduce the computational complexity (CC) per recovered symbol by recovering a multi-symbol output [10].

In this paper, for the first time, we present a workaround to solve the RNN parallelization issue, by using KD. The KD, in our case, is used to transfer the knowledge from a recurrent teacher model which is a bidirectional long-short term memory (biLSTM) coupled with 1D-convolutional NN (1D-CNN), the efficient equalizer model proposed in [11], to a feedforward student model: a 1D-CNN, shown in Fig. 1(a). Furthermore, we point out that together with the realization of a parallelizable structure, we reduce the inference latency of the model obtained via KD by 38%. Finally, we also observed a slight 0.5 dB reduction in Q-factor at optimum power.

2. Parallelization of Recurrent Equalizers via Knowledge Distillation

The recurrent nature of biLSTM, which takes into account the input of the current stage and the output of the previous stage, makes it very useful for learning sequential data. Furthermore, the input flowing in both directions further enhances the biLSTM learning. Due to the sequential processing of the double-recurrent setting and, consequently, limited parallelizability, the RNN-based model is not easily implementable for high-speed transmission. On the other hand, 1D-CNNs are feedforward-based NNs where the input temporal sequential batches are processed independently, hence, parallel operations are possible. Fig. 1(b) illustrates the recurrent structure at the top with a feedback loop, preventing parallelization, while the feedforward structure at the bottom can process multiple sets of inputs and provide multiple outputs simultaneously. Transforming the model architecture from the biLSTM to 1D-convolutional layers enables parallel computation for the previously proposed recurrent NN-based equalizer (NNE) or the teacher model.

Fig. 1(a) shows the teacher and student model structures and the KD process. The teacher NN contains a biLSTM layer with 100 hidden units and a linear 1D-CNN layer with two filters and a kernel size of 51. The student model has a dilated convolutional structure, which allows the NN to deal with long-term temporal dependencies and have large receptive fields within only a few layers [9, 12]. The last 1D-CNN layer of both the teacher and



Fig. 1: (a) KD framework with biLSTM+1D-CNN as a teacher model and dilated 1D-CNN as a student model; (b) Illustration of the parallelizability comparison between a recurrent cell (top) and a 1D-CNN (bottom).

the student has the same parameters. To mimic biLSTM, which learns the input data in forward and backward directions, the student model learns both directions of the training data. The backward direction input means the input sequence (forward direction) in reverse time order. The Bayesian Optimizer (BO) [3] is used to optimize the hyperparameter values of the student model. The estimated optimal values are depicted in Fig. 1(a). Note that (38, 23, 1) means that the 1D-CNN layer operates with 38 filters, a kernel size of 23, and a dilation rate of 1. The activation function of the dilated 1D-CNN part is LeakyRelu. The KD loss approach (Eq. (1)) for the student model training implies a joint loss function that takes into account both the loss between the student's and teacher's predictions and the loss between the student's predictions and "ground truth":

$$L_{\rm KD} = \alpha L_1(\hat{y}_S, \hat{y}_{\rm T}) + (1 - \alpha) L_2(\hat{y}_S, y_{\rm true}), \tag{1}$$

where y_{true} is the ground-truth labels, \hat{y}_s and \hat{y}_T represent the student's and the teacher's predictions, respectively, and α is the hyperparameter balancing the contribution of each term to the final loss. The teacher model is pretrained and used only to create teacher labels. The student model's training with KD is carried out for 1000 epochs, mini-batch size of 2¹⁰; the learning rate and α found by BO are equal to 0.00026 and 0.903, respectively. The dataset is obtained by numerical simulation, assuming the transmission of a single 30 GBd, 64-QAM dualpolarization channel along 20 × 50 km standard single mode fiber (SSMF) spans. At the end of each fiber span, optical fiber losses are compensated for by an Erbium-Doped Fiber Amplifier (EDFA) with a 4.5 dB noise figure.

Typically, the NNEs proposed in previous works are designed to recover one symbol at a time [3]. However, the single-symbol output NNEs can be computationally inefficient, as the weights and biases trained to recover one symbol may still be useful for recovering multiple symbols [10]. Therefore, we focus on multi-symbol output equalizers to reduce the CC per recovered symbol. The last layer of our model adopts the 1D-CNN layer containing two filters with a linear activation function to recover both real and imaginary parts of the signal. A set of 221 input symbols is fed to the NN simultaneously, to recover 171 symbols at the output at each inference step.

3. Results and Discussion

We compare our proposed student model (1D-CNN) using the KD framework with the teacher model (biLSTM +CNN) pre-trained as in [11], the student model trained from scratch (without KD) with exactly the same settings, and the student model trained from scratch with L2 regularizer [13]. The optimum L2 coefficient depends on the launch power. At 2 dBm launch power, the optimum L2 coefficient found by grid search is 10^{-4} . Fig. 2(a) depicts Q-factor vs. launch power for different types of NNEs. We can observe that the Q-factor performance of the feedforward student model with KD drops by 0.5 dB compared to the recurrent teacher model at its optimal launch power (2 dBm). With KD, the performance of the student model is comparable to that of the teacher model in the linear transmission regime, but the student's performance degrades slightly as the launch power increases. However, when training the student model from scratch without KD, the model suffers from overfitting, resulting in a noticeable degradation of the peak performance by 2.4 dB at its optimal power. Training the student model with the L2 regularizer, which helps enhance the generalization capability, improves the performance of the NN compared to training it from scratch only, but still does not reach a similar performance level as the one achieved when the student model is trained with KD. The performance achieved using digital back propagation (DBP) 1 step/span (STpS) and chromatic dispersion compensation (CDC) are also shown for reference.

Although the teacher model leads to slightly better optical performance than the student model with KD, Fig. 2(b) shows a reduction in the inference time per recovered symbol of the student model with KD of 17% and 38% when using the CPU and GPU, respectively. Note that both the teacher and the student models recover



Fig. 2: (a) Q-factor as a function of the launch power for the NNE obtained via KD, compared to the original (teacher) model, CDC, and 1 STpS DBP; (b) Inference time of teacher/student models; (c) Weight distribution of student model trained with different approaches.

171 symbols per inference step. Using a GPU leads to a more significant reduction in latency because of its better parallel-computing ability, which enables it to better exploit the parallelizability of the student feedforward structure. The parallelization of the proposed feedforward equalizer and its savings in latency are key to the real-time hardware implementation of NNEs.

Now, we study the features associated with the KD-trained model. For this purpose, we also report the weight distribution of the student model trained with different approaches in Fig. 2(c). Compared to the student model trained from scratch, the student model with KD has a more regularized weight distribution: the weights are more concentrated around zero. This characteristic helps reduce the model's variance and overfitting. The optimal value of α in the KD loss function is 0.903, which means that the student model learns 90.3% from the teacher labels and the rest comes from the ground-truth labels. This fact demonstrates the effectiveness of the teacher labels in the student's learning. The teacher constellation/labels depicted in Fig. 1(a) show that the teacher also provides helpful information on the noise, whereas this information cannot be encoded in the ground-truth labels (which contain only real values). The weight distribution of the student model with KD and the improvement in Q-factor, compared to the training of the 1D-CNN without KD, both support the concept of using teacher labels as efficient regularizers [8].

4. Conclusions

In this work, for the first time, the knowledge distillation technique has been proposed as an efficient tool to achieve the parallelizability of the recurrent equalizers: in our study, KD transfers the knowledge from a recurrent-connection-based biLSTM equalizer to a parallelizable feed-forward 1D-CNN. This approach enables the parallelization of signal processing, allowing us to essentially simplify the hardware implementation of NN models. We also show that the proposed feedforward equalizer obtained with KD, results in a reduced signal processing latency by 38% compared to the original biLSTM model, at the expense of slightly reducing the maximum Q-factor by 0.5 dB.

Acknowledgements: This work has received funding from the EU Horizon 2020 program under the Marie Skłodowska-Curie grant agreement No. 956713 (MENTOR) and 813144 (REAL-NET). SKT acknowledges the support of the EPSRC project TRANSNET (EP/R035342/1).

References

- T. Xu, et al., "Information rates in Kerr nonlinearity limited optical fiber communication systems", Opt. Express, vol. 29, no. 11, pp. 17428-17439, 2021.
- [2] J. C. Cartledge, et al., "Digital signal processing for fiber nonlinearities", Opt. Express, vol. 25, no. 3, pp. 1916-1936, 2017.
- [3] P. J. Freire, et al., "Performance Versus Complexity Study of Neural Network Equalizers in Coherent Optical Systems", J. Lightwave Technol., vol. 39, no. 19, pp. 6085-6096, 2021.
- [4] S. Deligiannidis, et al., "Compensation of fiber nonlinearities in digital coherent systems leveraging long short-term memory neural networks", J. Lightwave Technol., vol. 38, no. 21, pp. 5991-5999, 2020.
- [5] S. Deligiannidis, et al., "Performance and complexity analysis of bi-directional recurrent neural network models versus Volterra nonlinear equalizers in digital coherent systems", J. Lightwave Technol., vol. 39, no. 18, pp. 5791-5798, 2021.
- [6] A. X. M. Chang, et al., "Hardware accelerators for recurrent neural networks on FPGA", IEEE ISCAS, pp. 1-4, 2017.
- [7] R. Robey and Y. Zamora, "Parallel and high performance computing" (Simon and Schuster, 2021).
- [8] G. Hinton, et al., "Distilling the knowledge in a neural network", Deep Learning and Representation Learning Workshop, NIPS, 2014.
- Q. Xu, et al., "Contrastive adversarial knowledge distillation for deep model compression in time-series regression tasks", Neurocomputing, vol. 485, pp. 242-251, 2022.
- [10] B. Sang, et al., "Low Complexity Neural Network Equalization Based on Multi-Symbol Output Technique for 200+ Gbps IM/DD Short Reach Optical System", J. Lightwave Technol., vol. 40, no. 9, pp. 2890-2900, 2022.
- [11] P. J. Freire, et al., "Reducing Computational Complexity of Neural Networks in Optical Channel Equalization: From Concepts to Implementation", arXiv:2208.12866, 2022.
- [12] A. van den Oord, et al., "WaveNet: A Generative Model for Raw Audio", Proc. 9th ISCA Workshop on SSW, pp. 125, 2016.
- [13] A. Ng, et al., "Feature selection, L1 vs. L2 regularization, and rotational invariance", Proceedings of the Twenty-fFrst International Conference on Machine Learning, pp. 78, 2004.