



OPEN

Experimental implementation of a neural network optical channel equalizer in restricted hardware using pruning and quantization

Diego Argüello Ron^{1✉}, Pedro J. Freire^{1,2}, Jaroslaw E. Prilepsky¹, Morteza Kamalian-Kopae¹, Antonio Napoli² & Sergei K. Turitsyn^{1✉}

The deployment of artificial neural networks-based optical channel equalizers on edge-computing devices is critically important for the next generation of optical communication systems. However, this is still a highly challenging problem, mainly due to the computational complexity of the artificial neural networks (NNs) required for the efficient equalization of nonlinear optical channels with large dispersion-induced memory. To implement the NN-based optical channel equalizer in hardware, a substantial complexity reduction is needed, while we have to keep an acceptable performance level of the simplified NN model. In this work, we address the complexity reduction problem by applying pruning and quantization techniques to an NN-based optical channel equalizer. We use an exemplary NN architecture, the multi-layer perceptron (MLP), to mitigate the impairments for 30 GBd 1000 km transmission over a standard single-mode fiber, and demonstrate that it is feasible to reduce the equalizer's memory by up to 87.12%, and its complexity by up to 78.34%, without noticeable performance degradation. In addition to this, we accurately define the computational complexity of a compressed NN-based equalizer in the digital signal processing (DSP) sense. Further, we examine the impact of using hardware with different CPU and GPU features on the power consumption and latency for the compressed equalizer. We also verify the developed technique experimentally, by implementing the reduced NN equalizer on two standard edge-computing hardware units: Raspberry Pi 4 and Nvidia Jetson Nano, which are used to process the data generated via simulating the signal's propagation down the optical-fiber system.

Optical communications form the backbone of the global digital infrastructure. Nowadays, optical networks are the main providers of global data traffic, not only interconnecting billions of people, but also supporting the life-cycle of a huge number of different autonomous devices, machines, and control systems. One of the major factors limiting the throughput of contemporary fiber-optic communication systems is the nonlinearity-induced transmission impairments^{1,2}, emerging from both the fiber media's nonlinear response and the system's components. The existing and potential solutions to this problem include, e.g., the mid-span optical phase conjugation, digital back-propagation (DBP), and inverse Volterra series transfer function, to mention just a few noticeable methods²⁻⁴. But, it should be stressed that in the telecommunication industry, the competition between possible solutions occurs not only in terms of performance but also in terms of hardware deployment options, operational costs, and power consumption.

During the last years, the approaches based on machine learning techniques and, in particular, those utilizing NNs, have become an increasingly popular topic of research, as the NNs can efficiently unroll both fiber and component-induced impairments⁵⁻¹⁵. One of the straightforward ways of using an NN for signal's corruption compensation in optical transmission systems is to plug it into the system as a post-equalizer^{7,10,14}, a special signal processing device at the receiver side, aimed at counteracting the detrimental effects emerging during the data transmission¹⁶. Numerous preceding studies have demonstrated the potential of this type of solution^{7,8}. A number of NN architectures have already been analyzed in different types of optical systems (submarine, long-haul, metro, and access). These architectures include the feed-forward NN designs such as the MLP^{7,10,14,15}, considered in the current study, or more sophisticated recurrent-type NN structures^{10-12,17}. However, the practical deployment of

¹Aston Institute of Photonic Technologies, Aston University, Birmingham B4 7ET, UK. ²Infinera, Sankt-Martin-Str. 76, 81541 Munich, Germany. ✉email: d.arguelloron@aston.ac.uk; s.k.turitsyn@aston.ac.uk

real-time NN-based channel equalizers implies that their computational complexity is, at least, comparable, or, desirably lower than that of existing conventional digital signal processing (DSP) solutions¹⁸, and remains a matter of debate. This is a relevant aspect because the good performance achieved by the NNs is typically linked to the use of a large number of parameters and floating-point operations¹⁰. The high computational complexity leads, in turn, to high memory and computing power requirements, increasing the energy and resource consumption^{19,20}. Thus, the use of NN-based methods, while being, undoubtedly, promising and attractive, faces a major challenge in optical channel equalization, where the computational complexity emerges as an important limiting real-time deployment factor^{10,12,20,21}. We notice here that it is, of course, well known that some NN architectures can be simplified without significantly affecting their performance, thanks, e.g., to strategies such as pruning and quantization^{19,20,22–25}. However, their application in the experimental environment of the resource-restricted hardware has not been fully studied *in the context of coherent optical channel equalization* yet. It is also necessary to understand and further analyze the trade-off between the complexity reduction and the degradation of system performance, as well as the complexity reduction impact on the end device's energy consumption.

In this paper, we apply the pruning and quantization techniques to reduce the hardware requirements of an NN-based coherent optical channel equalizer, while keeping its performance at a high level. We also emphasize the importance of an accurate evaluation of the equalizer's computational complexity in the DSP sense. Apart from the complexity and inference time study, an additional novelty and advance of our work lies in the energy consumption analysis and the study of the impact that the characteristics of both the hardware and the model, have on these metrics.

Results

We develop and experimentally evaluate the performance of a low-complexity NN-based equalizer that can be deployed on resource-constrained hardware and, at the same time, can successfully mitigate nonlinear transmission impairments in a simulated optical communication system. This is achieved by applying the pruning and quantization techniques to the NN²³, and by studying the optimal trade-off between the complexity of the NN solution and its performance. The obtained results can be split into three main categories.

First, we quantify how complexity reduction techniques affect the performance of the NN model and establish a compression limit for optimal performance versus complexity trade-off. Second, we analyze the computational complexity of the pruned and quantized NN-based equalizer in terms of DSP. Finally, we experimentally evaluate the impact that the characteristics of the hardware and the NN model have on the signal processing time and energy consumption by deploying the latter on both a Raspberry Pi 4 and a Nvidia Jetson Nano.

Now we briefly review the previous results in the field of compression techniques applied to NN-based equalizers in optical links, to underline the novelty of our current approach. The use of these techniques to reduce the NNs complexity in optical systems is, clearly, not a new concept²⁵. However, the compression methods have recently gained a new wave of attention due to the question of how realistic the hardware implementation of NN-based equalizers in optical transmission systems is. In a direct detection transmission system, a parallel-pruned NN equalizer for a 100-Gbps PAM-4 links were tested experimentally using the enhanced version of the one-shot pruning method²⁶, which decreased by 50% the resource consumption without significant performance degradation. When considering coherent optical transmission, the complexity of the so-called learned DBP nonlinearity mitigation method was reduced by pruning the coefficients in the finite impulse response filters²⁷ (see more technical explanations in “[Methods](#)” section below). In that case, using a cascade of three filters, a sparsity level of around 92% can be achieved with a negligible impact on the overall performance. Recently, some advanced techniques for avoiding multiplications in such equalizers using additive powers-of-two quantization were tested²⁸. In the latter work, 99% of the weights could be removed using advanced pruning techniques, and instead of multiplications, just bit-shift operations were required. However, none of those works deal with the experimental demonstration of hardware implementation, and our study addresses exactly the latter problem.

So, unlike previous works, in the current study, we implement the compressed NN-based equalizer for the coherent optical channel in two different hardware platforms: a Raspberry Pi 4 and a Nvidia Jetson Nano. We also evaluate the impact of the compression techniques on the system's latency for each hardware type and study the performance-complexity trade-off. Finally, we carry out an analysis of energy consumption and of the impact that the characteristics of the hardware and the NN model have on it.

Optical communication system and equalizer design. To address the use of a MLP as an NN-based equalizer, an accurate measurement system for both the inference time and the power consumption, on both a Raspberry Pi and a Nvidia Jetson Nano, has been designed, so that the effects that pruning and quantization have on these metrics, can be characterized (see “[Methods](#)” section below for a detailed explanation). In Refs.^{10,14}, the non-compressed MLP post-equalizer was considered, and it was shown that it can successfully compensate for the nonlinearity-induced impairments in a coherent optical communication system. We analyze the equalizer's performance in terms of the standard achieved Q-factor, using the simulated data for a 0.1 root-raised cosine (RRC) dual-polarization signal, with 30 GBd, and 64-QAM modulation, for the transmission over the 20 × 50 km links of standard single-mode fiber (SSMF). We used the same simulator as described in Refs.^{10,29}, to generate our training and testing datasets, and the same procedure to training the NN-based equalizer (see “[Numerical setup and neural network model](#)” subsection in “[Methods](#)” for more details). In our configuration, the NN is placed at the receiver (Rx) side after the Integrated Coherent Receiver (ICR), Analog-Digital Converter (ADC), and DSP block. This last block consists of a matched filter and a linear equalizer. Concerning the matched filter, it is just the same RRC filter used in the transmitter. Moreover, the linear equalizer is composed of a full electronic chromatic dispersion compensation (CDC) stage and a normalization step, see Fig. 1. The CDC uses a frequency-domain equalizer and downsampling to the symbol rate, followed by a phase/amplitude normal-




Figure 1. Structure of a communication channel that is equalized using a pruned and quantized neural network deployed on resource-restricted hardware (e.g. a Raspberry Pi 4 or a Nvidia Jetson Nano).




Figure 2. Performance comparison for the NN-based equalizer with respect to the regular DSP.

izer to the transmitted ones. This normalization process can be viewed as its normalization by a constant \mathcal{K}_{DSP} learned using the following equation:

$$\mathcal{K}_{\text{DSP}} = \min_{\mathcal{K}} \|\mathcal{K} \times x_{h/v}(z, t) - x_{h/v}(0, t)\|, \quad (1)$$

where constants \mathcal{K} , $\mathcal{K}_{\text{DSP}} \in \mathbb{C}$ and $x_{h/v}$ is the signal in either h or v polarisation. No other distortions—related to the components within the transceiver—were considered.

For this system, the best optimal power occurred at -1 dBm with the Q-factor being close to 7.8, as it can be appreciated in Fig. 2. We then wanted to investigate the 3 next powers (e.g. 0 dBm, 1 dBm, and 2 dBm) going towards the higher nonlinear regime, where the task of the NN would be more complicated.

The hyperparameters that define the structure of the NN are obtained using a Bayesian optimizer (BO)^{10,30}, where the optimization is carried out with regards to the signal’s restoration quality performance (see “Numerical setup and neural network model” subsection in “Methods”). The resulting optimized MLP has three hidden layers (we did not optimize the number of layers, but the number of neurons and the activation functions type), with 500, 10, and 500 neurons, respectively. (These numbers were set as the minimal and maximal weights number limits, within which the BO algorithm was searching the optimal configuration). The activation function “tanh” was chosen by the optimizer and no bias is employed. The NN takes the downsampled signal (1 sample per symbol) and inputs into the equalizer $N = 10$ neighbors symbols (number of taps) to recover the central one. This memory size was defined by the BO procedure. The NN was subjected to pruning and quantization after it had been trained and tested. We analyzed the performance of different NN models depending on their sparsity level; the latter ranged from 20 to 90%, with a 10% increment. The weights and activations are quantized, converting their data type from 32-bit single-precision floating-point (FP32) to 8-bit integer (INT8). The quantization was carried out to enable a real-time use of the model as well as its deployment on resource-constrained hardware. The final system is depicted in Fig. 1. The inference process (the signal equalization) was, first, carried out using a MSI GP76 Leopard personal computer, equipped with Intel® Core™ i9-10870H processor, 32 GB of RAM and GPU Nvidia RTX2070. The results obtained on this computer were used as a benchmark and compared to those obtained on two small single-board computers: a Raspberry Pi 4 and a Nvidia Jetson Nano.

Finally, the NNs were developed using TensorFlow. The pruning and quantization techniques were implemented using the TensorFlow Model Optimization Toolkit—Pruning API and TensorFlow Lite³¹.

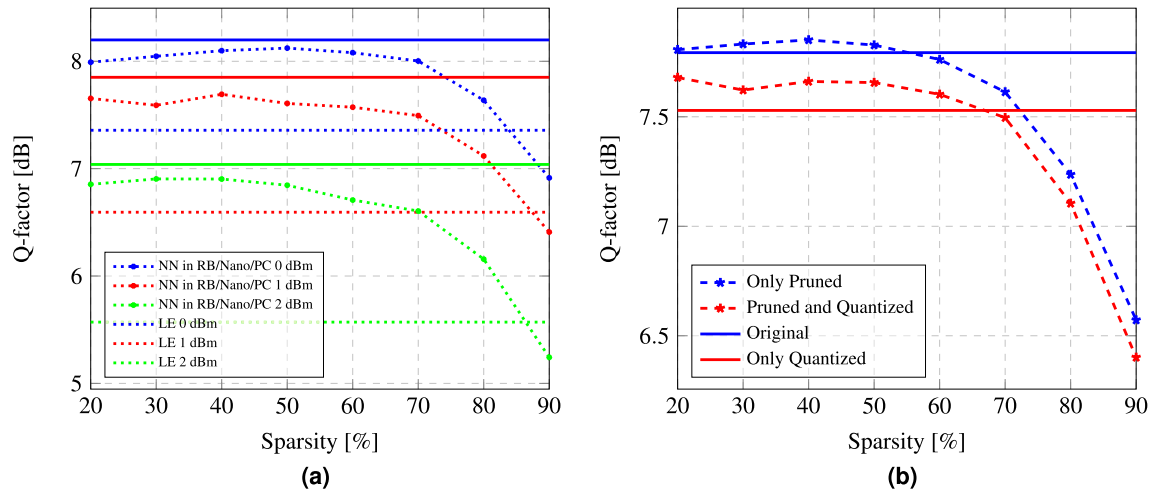


Figure 3. (a) Q-factor achieved for pruned and quantized models versus the level of sparsity for datasets corresponding to three launch powers: 0 dBm, 1 dBm, and 2 dBm; The solid lines correspond to the Q-factor achieved by the original model. The dashed lines show the Q-factor when only linear equalization (LE) is implemented. (b) Q-factor achieved after pruning compared to the one achieved after both pruning and quantization, for different levels of sparsity and for a dataset corresponding to the 1 dBm launch power. The blue and red solid lines correspond to the Q-factor achieved by the original model and the one achieved by this model after quantization, respectively.

Compressing process for neural network equalizers. When designing an NN for a particular purpose, the traditional approach consists in using dense and over-parametrized models, insofar as it often can provide a good model's performance and learning capabilities^{32,33}. This is due to the over-parametrization's smoothing effect on the loss function, which benefits the convergence of the gradient descent techniques used to optimize the model³². However, some precautions must be taken while training an over-parametrized model, because such models often tend to overfit, and their generalization capability can be degraded^{32,34}.

The good performance achieved due to over-parametrization comes at the cost of larger computational and memory resources. This also results in a longer inference time (latency growth) and higher energy consumption. Note that these costs are the consequence of parameter redundancy and a large number of floating-point operations^{20,23}. Therefore, the capabilities of high-complexity NN-based equalizers do not translate yet into end-user applications on resource-constrained hardware. Thus, reducing the gap between the algorithmic solutions and the experimental real-world implementations is an increasingly active topic of research. During the past several years, noticeable efforts have been invested in developing techniques that can help to simplify the NNs without significantly decreasing their performance. These techniques are grouped under the term "NNs compression methods", and the most common approaches are: down-sizing the models, factorizing the operators, quantization, parameter sharing or pruning^{20,23,24}. When these techniques are applied, the final model typically becomes much less complex, and, therefore, its latency, or the time it takes to make a prediction, decreases, which also results in a lower energy consumption²⁰. In this work, we focus on both pruning and quantization for compressing our NN equalizer and quantify a trade-off between complexity reduction and system performance, see "Methods" section for a detailed description of both approaches.

Performance vs. compression trade-off. Firstly, we note that the complexity reduction of the equalizer must not affect its performance drastically, i.e. the system's performance is still required to be within an acceptable range. In Fig. 3a, the Q-factor achieved by the NN equalizer is depicted versus different sparsity values, for three launch power levels: 0 dBm, blue; 1 dBm, red; and 2 dBm, green. The results are shown using dotted lines and stars, which are those obtained on the PC, Raspberry Pi, and Nvidia Jetson Nano, using the pruned and quantized model. For each of these launch powers, two baselines for the Q-factor are depicted: one corresponds to the level achieved by the uncompressed model, defined by the straight lines, while the other provides the benchmark when we do not employ any NN equalization and use only standard linear chromatic dispersion compensation plus phase/amplitude normalization (LE, linear equalization); the latter levels for the three different launch powers are marked by dotted lines having the appropriate colors.

Figure 3b quantifies the impact that each compression technique has on the performance: in that figure, we plotted the Q-factor achieved by the NN equalizer versus different values of sparsity, for the 1 dBm launch power. The blue and red straight lines represent the Q-factor of the original model and the Q-factor achieved by it after being quantized. The dotted lines with asterisks, show the performance of a model that has been only pruned (blue), and the performance in the case of both pruning and quantization (red). It is seen that a substantial reduction of the complexity can be achieved without a dramatic degradation of the performance. The sparsity levels at which the fast deterioration of the performance occurs, are also clearly seen in this figure.




Figure 4. Complexity and size reduction achieved via pruning and quantization for different levels of sparsity. The dashed black line represents the reference complexity when only quantization is applied.

First, it can be observed from Fig. 3a that the quantization and pruning process does not cause a significant performance degradation until a sparsity level equal to 60% is reached, with just a 4% performance reduction. However, when we move to sparsity levels around 90%, the performance is close to the one achieved using a linear equalization (i.e., the Q-factor curves drop to the levels marked with the dashed lines of the same color).

We can conclude that when the levels of sparsity are above 60%, the decrease in the performance is mainly the effect of the quantization process. A nearly 2.5% drop in the Q-factor value has also been observed when quantizing an already pruned model. Once the levels of sparsity are higher than 60%, the reduction in performance due to the quantization gets accelerated. Moreover, we observe that some degree of sparsification can even improve the model's performance with respect to the unpruned model. This behavior has already been reported in other studies and it was found that it is specifically pertinent to the over-parametrized models. Thus, the NNs with less complex structures do not show up such an increase in performance due to low-sparsity pruning, making it impossible to achieve such a good performance-complexity ratios^{32,33,35,36}.

Computational complexity analysis. Figure 4 depicts the reduction in the size of the model as well as the model's computational complexity for different sparsity values, after having applied quantization. For the definition of the metrics used to calculate the computational complexity as well as the size of the models, see the subsections "Computational complexity metrics and memory size metrics" in "Methods". Overall, we have achieved an 87.12% reduction in the memory size after pruning 60% of the NN equalizer weights and quantizing the remaining ones. As a consequence, the size of the model went down from 201.4 to 25.9 kilobytes. For the decrease of the model's computational complexity, it goes from 75,960,427.38 to 16,447,962 bit operations (BoPs) after applying the same compression strategy, which is a 78.34% reduction (see the explicit definition of BoPs in "Methods" section). We would like to point out once more that sparsity levels of 60% can be reached without a substantial performance loss. Therefore, approximately the same high level of performance can be achieved with a model that is significantly less complex than the initial NN structure, which is one of the main findings of our work.

It is worth mentioning the individual impact that quantization and pruning have on the computational complexity of the model. When the computational complexity is calculated for a quantized, but unpruned model, the number of BOPs is equal to 23,321,563. Therefore, if this value is compared with the already mentioned 75,960,427 BOPs for the unpruned and unquantized NN, a reduction in complexity of a 69.3% is obtained thanks to quantization. As it can be seen in Fig. 4, the remaining gain comes from the pruning technique, and it grows linearly as indicated in Eq. (5).

Online latency evaluation. Numerous deep learning applications are latency-critical, and therefore the inference time must be within the bounds specified by service level objectives. Optical communication applications that employ deep learning techniques are a good example of this. Note that the latency is highly dependent on the NN model implementation and the hardware employed (e.g., FPGA, CPU, GPU). Please refer to "Methods" section for more details on the devices' inference time measurements.

When measuring the inference time for the different types of hardware and the quantized model that has had 60% of its weights pruned, the results are:

- Latency Raspberry Pi: $\mu = 0.81$ s and $\sigma = \pm 0.035$
- Latency Nvidia Jetson Nano: $\mu = 0.53$ s and $\sigma = \pm 0.022$
- Latency PC: $\mu = 0.1$ s and $\sigma = 0.006$

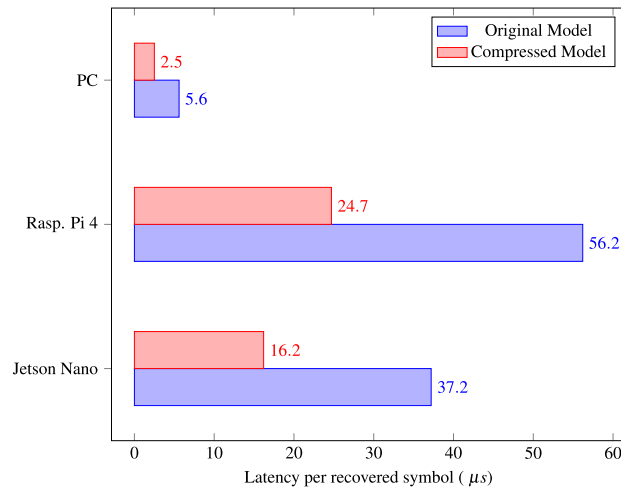


Figure 5. Summary of the symbol processing (inference) time for the compressed NN models (after pruning and quantization) and the original models for three devices under evaluation: a Raspberry Pi 4, a Nvidia Jetson Nano, and a standard PC.

In the case of the unpruned and unquantized model:

- Latency Raspberry Pi : $\mu = 1.84$ s and $\sigma = \pm 0.08$
- Latency Nvidia Jetson Nano: $\mu = 1.22$ s and $\sigma = \pm 0.052$ s
- Latency PC: $\mu = 0.18$ s and $\sigma = \pm 0.008$

Figure 5 shows the latency of the considered NN model before and after quantization. We notice that the results are expressed in a way that is more appropriate for the task at hand. Thus, latency is defined as the time it takes to process one symbol: we have averaged it over 30 k symbols. With the quantized model, we observe approximately a 56% reduction in latency for all three values of power, when compared to the original model. We must notice that pruning is not taken into account because it does not affect this metric since Tensorflow Lite does not support sparse inference yet, which makes the algorithm still use the same amount of cache memory. Also, we could observe that Raspberry Pi has the longest inference time among our devices. This is in line with the fact that Raspberry is designed as a low-cost and general-purpose single-board computer³⁷. On the other hand, the Nvidia Jetson Nano was developed with GPU capabilities, which makes it more suitable for deep learning applications, allowing us to achieve lower latencies.

Online energy consumption evaluation. Within the context of edge computing, not only is speed an important factor, but also power efficiency. In this work, the metric used to evaluate the energy consumption and compare the different types of hardware for the coherent optical channel equalization task is the energy per recovered symbol. When using a quantized model with a pruning level of 60%, the average energy consumed during inference for the Raspberry Pi 4 and the Nvidia Jetson Nano is 2.98 W ($\sigma = \pm 0.012$) and 3.03 W ($\sigma = \pm 0.017$), respectively. On the other hand, if the original model is employed, there is an increase in energy consumption of around 3%, which is congruent with the findings in previous works²³. Thus, during inference, the Raspberry Pi 4 consumes 3.06 W ($\sigma = \pm 0.011$) and the Nvidia Jetson Nano 3.13 W ($\sigma = \pm 0.015$), respectively. Multiplying these values by the NN processing times per recovered symbol reported in Fig. 5, we obtain the results presented in Fig. 6. We note that Raspberry Pi has the highest energy consumption per recovered symbol. This is a consequence of the lack of a GPU, which results in longer inference times. Thus, the Nvidia Jetson Nano consumes 33.78% less energy than the Raspberry Pi 4. Regarding pruning and quantization, the use of these techniques allows an energy saving of 56.98% for the Raspberry Pi 4 and a 57.76% saving for the Nvidia Jetson Nano.

It must be noticed that although TensorFlow Lite does not support sparse inference and therefore pruning does not help to reduce the inference time, it affects the size of the model. This has a direct effect on the power consumption of the device due to the decrease in the use of resources. In contrast, quantization has a positive effect on both of these parameters thanks to employing lower precision formats and reducing the size of the model. Therefore, it has a stronger effect on energy consumption. This is reflected in the results exposed in this section. Moreover, it is congruent with the findings reported in previous studies^{23,38}.

See “Methods” section for more details on the energy consumption measurement.

Discussion

In our work, we investigated how we can use pruning and quantization to reduce the complexity of the hardware implementation of an NN-based channel equalizer in a coherent optical transmission system. With this, we tested the implementation of the designed equalizer experimentally, using a Raspberry Pi 4 and a Nvidia Jetson

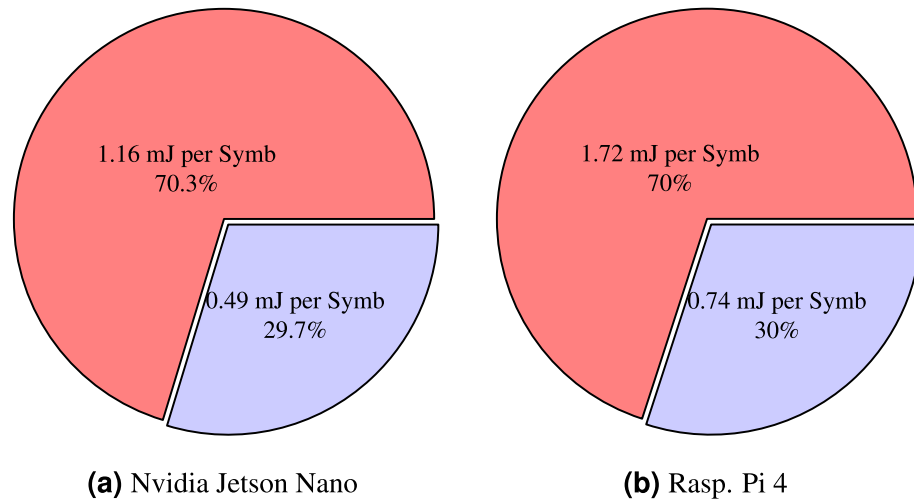


Figure 6. Energy consumption for Raspberry Pi 4 and Nvidia Jetson Nano. The blue section represents the energy consumption per recovered symbol when using the compressed model, and its relative energy cost is expressed as a percentage with respect to the sum of the energy consumed by both the original and compressed models. Likewise, the red section describes the energy consumption per recovered symbol when using the original model and its relative energy cost.

Nano. It was demonstrated that it is possible to reduce the NN's memory usage by 87.12%, and the NN's computational complexity by 78.34% without any serious penalty in performance, thanks to the two aforementioned compression techniques.

Moreover, the effect of using different types of hardware was experimentally characterized by measuring the inference time and energy consumption in both a Raspberry Pi 4 and a Nvidia Jetson Nano. We note, however, that we experimented only with the edge devices, and the data from the communication system were obtained via simulations; but we do not expect that the results regarding the performance vs complexity trade-off achieved thanks to pruning and quantization for the true optical system would seriously differ. It has been demonstrated that the Nvidia Jetson Nano allows 34% faster inference times than the Raspberry Pi, and that, thanks to the quantization process, a 56% inference time reduction can be achieved. Finally, due to the use of pruning and quantization techniques, we achieve 56.98% energy savings for the Raspberry Pi 4 and 57.76% for the Nvidia Jetson Nano; we also observed that the latter device consumes 33.78% less energy.

Overall, our findings demonstrate that the usage of pruning and quantization can be a suitable strategy for the implementation of NN-based equalizers that are efficient in high-speed optical transmission systems when deployed on resource-restricted hardware. We believe that these model compression techniques can be used for the deployment of NN-based equalizers in real optical communication systems, and for the development of novel online optical signal processing tools. We hope that our results can also be of interest to the researchers developing sensing and laser systems, where the application of machine learning for field processing and characterization is a rapidly developing area of research³⁹.

Methods

Numerical setup and neural network model. We numerically simulated the dual-polarization (DP) transmission of a single-channel signal at 30 GBd. The signal is pre-shaped with a root-raised cosine (RRC) filter with 0.1 roll-off at a sampling rate of 8 samples per symbol. In addition, the signal modulation format is 64-QAM. We considered the case of transmission over 20×50 km links of SMF. The optical signal propagation along the fiber was simulated by solving the Manakov equation via split-step Fourier method⁴⁰ with the resolution of 1 km per step. The considered parameters of the TWC fiber are: the attenuation parameter $\alpha = 0.23 \text{ dB/km}$, the dispersion coefficient $D = 2.8 \text{ ps}/(\text{nm} \times \text{km})$, and the effective nonlinearity coefficient $\gamma = 2.5 \text{ (W} \times \text{km)}^{-1}$. The SSMF parameters are: $\alpha = 0.2 \text{ dB/km}$, $D = 17 \text{ ps}/(\text{nm} \times \text{km})$, and $\gamma = 1.2 \text{ (W} \times \text{km)}^{-1}$. Moreover, after each span, an optical amplifier with the noise figure $\text{NF} = 4.5 \text{ dB}$ was placed to fully compensate fiber losses and added amplified spontaneous emission (ASE) noise. At the receiver, a standard Rx-DSP was employed. It consisted of the full electronic chromatic dispersion compensation (CDC) using a frequency-domain equalizer, the application of a matched filter, and the downsampling to the symbol rate. Finally, the received symbols were normalized (by phase and amplitude) to the transmitted ones. In this work, no additional transmitter distortions were taken into account. After the Rx-DSP, the bit error rate (BER) is estimated using the transmitted symbols, received soft symbols, and hard decisions after equalization.

The NN receives as input a tensor with a shape defined by three dimensions: $(B, M, 4)$, where B is the mini-batch size, M is the memory size determined by the number of neighbors N as $M = 2N + 1$, and 4 is the number of features for each symbol, which correspond to the real and imaginary parts of two polarization components. The NN will have to recover the real and imaginary parts of the k -th symbol of one of the polarization. Therefore the shape of the NN output batch can be expressed as $(B, 2)$. This task can be treated as a regression or

classification one. This aspect has been considered in previous studies and stated that the results achieved by regression and classification algorithms are similar but fewer epochs are needed in the case of regression. Thus, the mean square error (MSE) loss estimator is used in this paper, as it is the standard loss function employed in regression tasks⁴¹. The loss function is optimized using the Adam algorithm⁴² with the default learning rate equal to 0.001. The maximum number of epochs during the training process was 1000, as it was stopped earlier if the value of the loss function did not change over 150 epochs. After every training epoch, we calculated the BER obtained using the testing dataset. The optimal number of neurons and activation functions in each layer of the NN, as well as the memory (input) of the system were inferred employing the Bayesian Optimization algorithm (BO). The values tested for the number of neurons were $n \in [10, 500]$. For the activation function, the BO had to choose between: “tanh”, “ReLU”, “sigmoid” and “LeakyReLU”. The values tested for the memory (input) of the system were $N \in [5, 50]$. The metric of the BO was the BER, finding the hyperparameters that helped to reduce the BER as much as possible with a validation dataset of 2^{17} data points. The final solution was the use of “tanh” as an activation function and 500, 10, and 500 neurons for the first, second, and third layer, respectively. The training and test datasets were composed of independently generated symbols of length 2^{18} each. To prevent any possible data periodicity and overestimation^{43,44}, a pseudo-random bit sequence (PRBS) of order 32 was used to generate those datasets with different random seeds for each of them. The periodicity of the data is, therefore, 2^{12} times higher than our training dataset size. For the simulation, the Mersenne twister generator⁴⁵ was used with different random seeds. Moreover, the training data was shuffled before being used as an input to the NN.

Finally, we would like to notice an important matter as it is the necessity of the periodical retraining of the equalizer on realistic transmission. In this case, it may be a point of concern. This issue has already been addressed in previous studies²⁹, where it has been demonstrated that using transfer learning can drastically reduce the training time and training data requirements when changes on the transmission setup occur.

Pruning technique. With pruning, the redundant NN elements can be removed to sparsify the network without significantly limiting its ability to carry out a required task^{24,32,46}. Thus, networks with a reduced size and computational complexity are obtained, resulting in lower hardware requirements as well as faster prediction times^{23,24}. Furthermore, pruning acts as a regularization technique, improving the model quality by helping to reduce overfitting³². Moreover, retraining an already pruned NN can help to escape local loss function minima, which can lead to a better prediction accuracy²⁴. Thus, less complex models can often be achieved without a noticeable effect on the NN’s performance³².

Depending on what is going to be pruned, the sparsification techniques can be classified into two types: model sparsification and ephemeral sparsification³². In the first case, the sparsification is permanently applied to the model, while in the second case, the sparsification only takes place during the computing process. In our work, we will use the model sparsification, because of the effects it has on the final NN’s computing and memory hardware requirements. Adding to this, the model sparsification can consist in removing not only weights but also larger building blocks, such as neurons, convolutional filters, etc.³². Here we apply pruning to just the weights of the network, for the sake of simplicity and as far as it matches the NN structure (the MLP) that is considered.

After having defined what to prune, it is necessary to define when the pruning occurs. Based on this, there are two main types of pruning: static and dynamic²⁴. In the static case, the elements are removed from the NN after the training, and in this work, to demonstrate the effect, we use the static pruning variant because of its simplicity.

The static pruning is generally carried out in three steps. First, we decide upon what requires to be pruned. A simple approach to define the pruning objects can be to evaluate the NN’s performance with and without particular (pruned) elements. However, this poses scalability problems: we have to evaluate the performance when pruning each particular NN’s parameters, and there may be millions of these.

Alternatively, it is possible to select the elements to be removed randomly, which can be done faster^{32,47,48}. Following this latter approach, we beforehand decided to prune the weights. Once it has been decided which elements are to be pruned, it is necessary to establish the criteria for how the elements are to be removed from the NN, ensuring that high levels of sparsity are achieved without a significant loss in performance. When pruning the weights of the NN, it is possible to remove them based on different aspects: considering their magnitude (i.e., the weights having values close to zero are to be pruned, with the pruning percentage is defined by the sparsity level we aim to achieve), or their similarity (if two weights have a similar value, only one of those is kept); we mention that the other selection procedures also exist^{32,48}. Here, we pick the relatively simple weights pruning strategy based on their magnitude. In Fig. 7 we show the impact when we have pruned our NN equalizer by 40%. When comparing the weight distributions of the original and pruned models, it is clear that the sparsity level defines the number of weights that need to be pruned. Thus, the pruning process starts by removing the smallest weight and continues until the desired sparsity level is reached. Finally, a retraining or fine-tuning phase should be done, to reduce the degradation in the modified NN performance²⁴.

When carrying out pruning using the Tensorflow Model Optimization API, it is necessary to define a pruning Schedule to control this process by notifying at each step the level at which the layer should be pruned⁴⁹. In this work, the schedule known as Polynomial Decay is employed. The main characteristic of this type of schedule is that a polynomial sparsity function is built. In this case, the power of the function is equal to 3 and the pruning takes place every 50 steps. This means that during the last steps higher ratios of sparsification are employed (e.g. more weights are removed), speeding up the pruning process. On the other hand, if the power of the function were negative, pruning would be slowed-down. The model starts with a 0% sparsity and the process takes place during 300 epochs. This is approximately 35 % of the number of iterations required for training the original model. It is the objective of future works to optimize the hyperparameters of the pruning process, improve its efficiency and reduce the cost related to a high number of iterations.

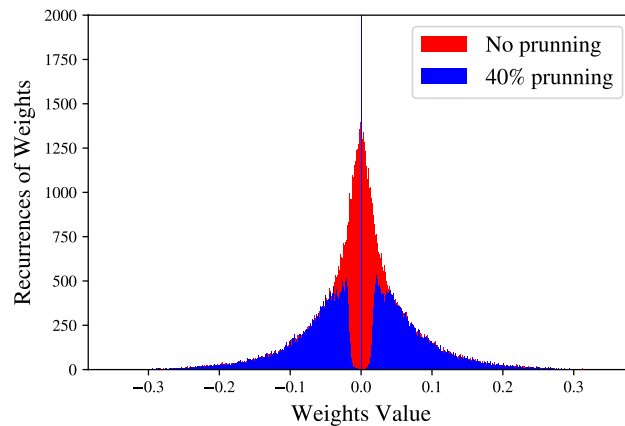


Figure 7. A typical distribution of the weights of the NN-based MLP equalizer without pruning and with pruning when the sparsity level is set to 40%.

Quantization technique. Besides the reduction in the number of operations involved in the NN signal processing, the precision of such arithmetic operations is another crucial factor when determining the model's complexity and, therefore, the inference latency, as well as equalizer's memory and energy requirements^{23,50–52}. The process of approximating a continuous variable with a specified set of discrete values is known as quantization. The number of discrete values will determine the number of bits necessary to represent the data. Thus, when applying this technique in the context of deep learning, the objective is decreasing the numeric precision used to encode the weights and activations of the models, avoiding a noticeable decrease in the NN's performance^{20,52}.

Using low-precision formats allows us to speed up math-intensive operations, such as convolution and matrix multiplication⁵². On the other hand, the inference (signal processing) time depends not only on the format representation of the digits involved in the mathematical operations but is also affected by transporting the data from memory to the computing elements^{23,38}. Moreover, heat is generated during the latter process and, therefore, using a lower-precision representation can result in energy savings²³. Finally, another benefit of using low-precision formats is that a reduced number of bits is needed to store the data, which reduces the memory footprint and size requirements^{23,52}.

FP32 has been traditionally used as the numerical format for encoding weights and activations (output of the neurons) in an NN, to take advantage of a wider dynamic range. However, as it has already been mentioned, this results in higher inference times, which is an important factor when a real-time signal processing is considered²⁰. A variety of alternatives to the FP32 numerical format for NN's elements representation have been proposed lately, to reduce the inference time, as well as to decrease the hardware requirements. For example, it is becoming popular to train NNs in FP16 formats, as it is supported by most deep learning accelerators²⁰. On the other hand, math-intensive tensor operations executed on INT8 types can see up to a 16× speed-up compared to the same operations in FP32. Moreover, memory-limited operations could see up to a 4× speed-up compared to the FP32 version^{22–24,52}. Therefore, in addition to pruning, we will reduce the precision of the weights and activations to further decrease the computational complexity of the equalizer, employing the technique known as integer quantization⁵².

The integer quantization maps a floating point value $x \in [\alpha, \beta]$ to a bit integer $x_q \in [\alpha_q, \beta_q]$. This mapping can be defined mathematically using the following formula: $x_q = \text{round}(\frac{1}{s}x + z)$, where s (a positive floating point number) is known as the scale, and z is the zero point (an integer). The scaling factor basically divides a range of real values, in this case those within the clipping range $[\alpha, \beta]$, into a number of partitions. Thus, it can be expressed as $s = \frac{\beta - \alpha}{2^b - 1}$ where b is the quantization bit width. On the other hand, the zero point can be defined as $z = \frac{\alpha(1 - 2^b)}{\beta - \alpha}$. Therefore, it will be 0 in the case of symmetric quantization. Moreover, the previous mapping can be refactored in order to take into account that if x is outside of the range $[\alpha, \beta]$, then x_q is outside of $[\alpha_q, \beta_q]$. Thus, it is necessary to clip the values when this happens; as a consequence, the mapping formula becomes: $x_q = \text{clip}(\text{round}[\frac{1}{s}x + z], \alpha_q, \beta_q)$, where the clip function takes the values^{24,53}:

$$\text{clip}(x, l, u) = \begin{cases} l & \text{if } x < l, \\ x & \text{if } l \leq x \leq u, \\ u & \text{if } x > u. \end{cases}$$

Integer quantization can take different forms, depending on the spacing between quantization levels and the symmetry of the clipping range (determined by the value of the zero-point z)⁵³. For the sake of simplicity, in this work, we used symmetric and uniform integer quantization.

The quantization process can occur after the training or during it. The first case is known as post-training quantization (PTQ) and the second one is the quantization aware training^{22–24}. In PTQ, a trained model has its weight and activations quantified. After this, a small unlabelled calibration set is used to determine the activations' dynamic ranges^{23,52–54}. No retraining is needed, which makes this method very popular because of its simplicity and lower data requirements^{53,54}. Nonetheless, when a trained model is directly quantized, this may

perturb the trained parameters, moving the model away from the convergence point reached during the training with a floating-point precision. In other words, we notice that PTQ can have accuracy-related issues⁵³.

In this work, the quantization is carried out after the training stage, i.e., we use the PTQ. The calibration process required to estimate the range, i.e. (min, max) of the activations in the model, is done by running a few inferences with a small portion of the test dataset. In our case, it consisted of 100 samples. When using the Tensorflow Lite API, the calibration is carried out automatically, and it is not possible to choose the number of inferences.

Computational complexity metrics. Finally, it is important to discuss how we can correctly evaluate the computational complexity of such models. In this regard, we quantitatively evaluate the reduction of computation complexity achieved by applying pruning and quantization, calculating the number of bits used during an inference step. The most common operations in an NN are multiply-and-accumulate operations (MACs). These are operations of the form $a = a + w \times x$, where three terms are involved: firstly, x corresponds to the input signal of the neuron; secondly, w refers to the weight; and, finally, the accumulate variable a ⁵⁵. Traditionally, the network complexity arithmetic has been measured using the number of MAC operations. However, in terms of the DSP processing, the number of BoPs is a more appropriate metric to describe the computational complexity of the model, as for low-precision networks composed of integer operations, it is not possible to measure the computational complexity using FLOPS^{22,56}. Thus, in this work, we use BoPs to quantify the complexity of the equalizer. It is important to notice that within the context of optical channel non-linear compensation, the complexity of NN-based channel equalizers has been traditionally measured taking into account only the number of multiplications^{12,44,57}. Thus, the accumulator contribution was neglected. However, in this project, we aim to have a more general complexity metric and therefore include it in our calculations.

The BOPs measure was proposed for the first time in⁵⁶, and defined for a convolutional layer that had been quantized as:

$$\text{BoPs} = mnk^2(b_a b_w + b_a + b_w + \log_2(nk^2)). \quad (2)$$

In Eq. (2), b_w and b_a are the weight and activation bit-width, respectively; n is the number of input channels, m is the number of output channels, and k defines the filters size (e.g. $k \times k$ filters)⁵⁸. Taking into account that a MAC operation takes the form: $a = a + w \times x$, it is possible to distinguish two contributions in the equation above: one corresponding to the $nk^2 \times b_0$ number of additions, where $b_0 = b_a + b_w + \log_2(nk^2)$ (e.g. accumulator width in the MAC operations), and the other corresponds to the number of multiplications, e.g. $nk^2(b_a b_w)$ ⁵⁶.

Equation (2) was further adapted for the case of a dense layer that has been both pruned and quantized⁵⁹. Thus, it is applicable to our case, as the MLP consists of a series of dense layers arranged one after the other:

$$\text{BoPs}_i = m_i n_i [(1 - f_{p_i}) b_{a_i} b_{w_i} + b_{a_i} + b_{w_i} + \log_2(n_i)], \quad (3)$$

In Eq. (3), n and m correspond to the number of inputs and outputs, respectively; b_w and b_a are the bit widths of the weights and activations. The additional term, f_{p_i} , is the fraction of pruned layer weights, which allows us to take into account the reduction in multiplication operations because of pruning. This is the reason why it only relates to the term $b_a b_w$ ⁵⁹.

Therefore, in our case of the MLP with 3 hidden layers, the total number of BOPs is:

$$\text{BoPs} = \text{BoPs}_{\text{input}} + \sum_i \text{BoPs}_i + \text{BoPs}_{\text{output}}, \quad (4)$$

where $i \in [1, 2, 3]$, $\text{BoPs}_{\text{input}}$ and $\text{BoPs}_{\text{output}}$ correspond to the contributions of the input and output layers. Equation (4) can be written in a less compact way as follows:

$$\begin{aligned} \text{BoPs}_{\text{MLP}} = & (n_i n_1 b_i + n_1 n_2 b_a + n_2 n_3 b_a + n_3 n_o b_a)(1 - f_p) b_w + (n_i n_1)(b_i + b_w) \log_2(n_i) \\ & + (n_1 n_2)(b_a + b_w) \log_2(n_1) + (n_2 n_3)(b_a + b_w) \log_2(n_2) + (n_3 n_o)(b_a + b_w) \log_2(n_3), \end{aligned} \quad (5)$$

where n_i , n_1 , n_2 , n_3 , and n_o are the number of neurons in the input, first, second, third, and output layers, respectively; b_w , b_a , b_o and b_i are the bit width of the weights, activations, output and input, respectively; f_p is the fraction of the weights that have been pruned in a layer, which, in our case, is the same for every layer.

Memory size metrics. In this work, the size of the model is defined as the number of bytes that it occupies in memory. Moreover, we notice the direct correlation between the value of this metric and the format used to represent the model. Thus, in contrast to the traditional formats used in Tensorflow (e.g. .h5 or HDF5 binary data format and .pb or protobuf), a TensorFlow Lite model is represented in a special efficient portable format identified by the .tflite file extension. This provides two main advantages: a reduced model's size and lower inference times. Therefore, the deployment of the NN model on a resource-restricted hardware becomes feasible. As a consequence, it would not make sense to compare the models saved in the traditional Tensorflow format with those that have been pruned and quantized as well as converted into Tensorflow Lite. We were aware of this situation during the realization of the procedure and, thus, to avoid overestimating the benefits of pruning and quantization, the unpruned and unquantized model were converted to .tflite format. To better understand the implications that this step has, the size of the original model in .h5 format would experiment a 96.22% size reduction after being converted to .tflite format, quantized and pruned (60% sparsity). On the other hand, if the original model has already been converted to .tflite, the size reduction is 87.12%. Of course, based on this, always using .tflite format instead of the other conventional ones seems to be the best strategy. The main reason behind not doing this is that a graph that is in .tflite format can not be trained again, as it only sup-

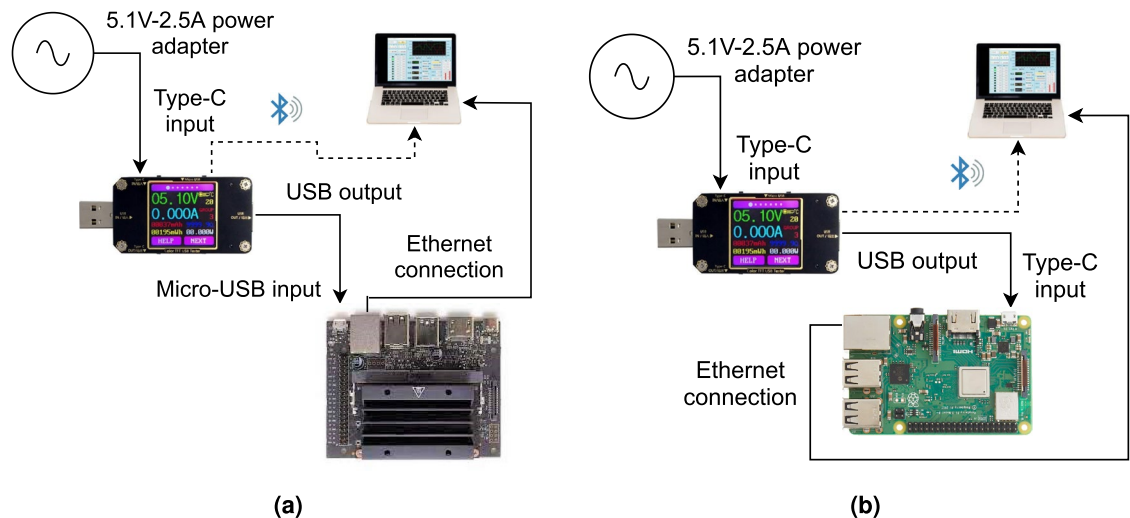


Figure 8. (a) The power measurement set-up for Nvidia Jetson Nano, and (b)—the same for Raspberry Pi.

ports an online inference mode. Nevertheless, a model that is, for example, in .h5 format, can be trained offline. Therefore, the .tfLite is only intended to be used in the context of edge computing.

Memory and processor restricted hardware. In many deep learning applications, low power consumption and a reduced inference time are especially desirable. Moreover, the use of graphics processing units (GPU) to attain high performance has some costs-related issues which are far from being ultimately solved^{37,60}. Therefore, a small, portable, and low-cost hardware is required to bring the solution to this problem. As a result, single-board computers have become popular, and Raspberry Pi 4 and Nvidia Jetson Nano are among the most used ones³⁷. Hence, here we analyse the functioning of our NN-based equalizer using these two aforementioned popular hardware types.

Raspberry Pi. Raspberry Pi is a small single-board computer. It is equipped with a Broadcom Video Core VI (32-bit) GPU, Quad-core ARM CortexA72 64-bit 1.5 GHz CPU, 2 USB 2.0 ports, and 2 USB 3.0 ports; for data storage, it uses a MicroSD card. Moreover, connections are provided through a Gigabit Ethernet/WiFi 802.11ac. It uses an OS known as Raspbian and has no GPU capability as well as no specialized hardware accelerator^{37,61}.

Nvidia Jetson Nano. Nvidia Jetson Nano is a small GPU-based single-board computer that allows the parallel operation of multiple NNs. It has a reduced size (100 mm × 80 mm × 29 mm) and is equipped with a Maxwell 128-core GPU, Quad-core ARM A57 64-bit 1.4 GHz CPU. Like in the case of Raspberry Pi, a MicroSD card is used to store the data. Finally, connections are established via Gigabit Ethernet and the OS employed is Linux-4Tegra, based on Ubuntu 18.04^{37,60}.

Power measurement. In this work, together with the latency and accuracy attributed to each model processing, we also address the issue of the power consumption for the NN equalizers implemented in the Nvidia Jetson Nano and the Raspberry Pi 4.

It is possible to measure the power consumption of both the Nvidia Jetson Nano and the Raspberry Pi in different ways. Regarding Nvidia Jetson Nano, there are three onboard sensors located at the power input, at the GPU, and at the CPU. Thus, the precision of the measurements is limited by these sensors. To read the recordings of these sensors, it is possible to do it automatically using the `tegrastats` tool, or manually by reading `.sys` files, a pseudo-file system on Linux. By using both approaches, the information of measurements for the power, voltage, and current can be readily collected⁶². In contrast, Raspberry Pi 4 has no system to easily provide power consumption numbers. Some software-based methods have been developed, as well as some empirical estimations⁶³. However, it has been demonstrated that most of the aforementioned software methods give just an approximation that may not be used if very precise results are required⁶³. On the other hand, the second empiric strategy to measure the power consumption on Raspberry Pi is specific for this type of hardware and cannot be used in Nvidia Jetson Nano.

To compare the power consumption of the equalizer on these two types of hardware, it is more accurate and desirable to use the same method in both of them, to avoid any instrumental bias. In this paper, we developed a platform-agnostic method through the use of a digital USB multimeter. The proposed power consumption measurement system addresses the problem of these devices having no onboard shunt resistors; such an approach allows us to easily measure power with an external energy probe. A schematic of the measurement set-ups is given in Fig. 8.

In the case of Raspberry Pi, the power is supplied through a USB type C port via a 5.1 V–2.5 A power adapter. For Nvidia Jetson Nano, the power can be supplied through a Micro-USB connector using a 5.1 V–2.5 A power

adapter or a Barrel jack 5 V–4 A (20 W) power supplier. It is possible to change from one configuration to the other by setting a jumper and moving from the 5 W Mode to the 10 W one. To use the same source of power as in Raspberry Pi, the Micro-USB configuration is used.

As energy is supplied through a USB connection, it is possible to measure the power using a USB digital multimeter. The model used in this work is the A3-B/A3 manufactured by Innovateking-EU. It records voltage, current, impedance, and power consumption. The input voltage and current ranges are 4.5 V–24 V and 0 A–3 A, respectively. Moreover, we can measure the energy in a range that goes from 0 to 99,999 mWh. The voltage and current measurement resolution are 0.01 V and 0.001 A, with the measurement accuracies $\pm 0.2\%$ and $\pm 0.8\%$, respectively.

The USB digital multimeter A3-B/A3 comes with the software named UM24C PC Software V1.3, which allows sending the measured data to a computer in real-time, as it is shown in Fig. 8a,b. During the measurement process, no peripherals are connected either to Raspberry Pi or Nvidia Jetson Nano, except for the Ethernet port. This is used for communication over SSH, Fig. 8. Moreover, 25 measures were taken for each device. In each of them, 100 inferences were run, and the power consumption was averaged over them, not taking into account the power consumed during the initialization phase.

Inference time measurement. To evaluate the inference time for each model, no peripherals are connected either to the Raspberry Pi or to the Nvidia Jetson Nano, except the Ethernet port, which is used to establish communication over the Secure Shell protocol. Moreover, any initialization time (e.g., library loading, data generation, and model weight loading) is ignored because this is a one-time cost that occurs during the device's setup. Furthermore, 25 measures were taken for each device. In each of them, 100 inferences were run (in each inference 30 k symbols are recovered) and the inference time was averaged, not taking into account the initialization phase.

Data availability

Data underlying the results presented in this paper are not publicly available at this time, but can be obtained from the authors upon request.

Received: 6 January 2022; Accepted: 3 May 2022

Published online: 24 May 2022

References

- Winzer, P. J., Neilson, D. T. & Chraplyvy, A. R. Fiber-optic transmission and networking: The previous 20 and the next 20 years. *Opt. Express* **26**, 24190–24239. <https://doi.org/10.1364/OE.26.024190> (2018).
- Cartledge, J. C., Guiomar, F. P., Kschischang, F. R., Liga, G. & Yankov, M. P. Digital signal processing for fiber nonlinearities. *Opt. Express* **25**, 1916–1936. <https://doi.org/10.1364/OE.25.001916> (2017).
- Rafique, D. Fiber nonlinearity compensation: Commercial applications and complexity analysis. *J. Lightw. Technol.* **34**, 544–553. <https://doi.org/10.1109/JLT.2015.2461512> (2016).
- Dar, R. & Winzer, P. J. Nonlinear interference mitigation: Methods and potential gain. *J. Lightw. Technol.* **35**, 903–930. <https://doi.org/10.1109/JLT.2016.2646752> (2017).
- Musumeci, F. *et al.* An overview on application of machine learning techniques in optical networks. *IEEE Commun. Surv. Tutor.* **21**, 1383–1408. <https://doi.org/10.1109/COMST.2018.2880039> (2019).
- Nevin, J. W. *et al.* Machine learning for optical fiber communication systems: An introduction and overview. *APL Photon.* <https://doi.org/10.1063/5.0070838> (2021).
- Jarajreh, M. A. *et al.* Artificial neural network nonlinear equalizer for coherent optical ofdm. *IEEE Photon. Technol. Lett.* **27**, 387–390. <https://doi.org/10.1109/LPT.2014.2375960> (2015).
- Häger, C. & Pfister, H. D. Nonlinear interference mitigation via deep neural networks. In *2018 Optical Fiber Communications Conference and Exposition (OFC)*, 1–3 (IEEE) (2018).
- Zhang, S. *et al.* Field and lab experimental demonstration of nonlinear impairment compensation using neural networks. *Nat. Commun.* **10**, 3033. <https://doi.org/10.1038/s41467-019-10911-9> (2019).
- Freire, P. J. *et al.* Performance versus complexity study of neural network equalizers in coherent optical systems. *J. Lightw. Technol.* **39**, 6085–6096. <https://doi.org/10.1109/JLT.2021.3096286> (2021).
- Deligiannidis, S., Bogris, A., Mesaritakis, C. & Kopsinis, Y. Compensation of fiber nonlinearities in digital coherent systems leveraging long short-term memory neural networks. *J. Lightw. Technol.* **38**, 5991–5999. <https://doi.org/10.1109/JLT.2020.3007919> (2020).
- Deligiannidis, S., Mesaritakis, C. & Bogris, A. Performance and complexity analysis of bi-directional recurrent neural network models versus volterra nonlinear equalizers in digital coherent systems. *J. Lightw. Technol.* **39**, 5791–5798. <https://doi.org/10.1109/JLT.2021.3092415> (2021).
- Freire, P. J. *et al.* Experimental study of deep neural network equalizers performance in optical links. In *2021 Optical Fiber Communications Conference and Exhibition (OFC)*, 1–3 (2021).
- Sidelnikov, O., Redyuk, A. & Sygletos, S. Equalization performance and complexity analysis of dynamic deep neural networks in long haul transmission systems. *Opt. Express* **26**, 32765–32776. <https://doi.org/10.1364/OE.26.032765> (2018).
- Sidelnikov, O. S., Redyuk, A. A., Sygletos, S. & Fedoruk, M. P. Methods for compensation of nonlinear effects in multichannel data transfer systems based on dynamic neural networks. *Quantum Electron.* **49**, 1154. <https://doi.org/10.1070/QEL17158> (2019).
- Barry, J. R., Lee, E. A. & Messerschmitt, D. G. *Digital Communication* 3rd edn. (Springer, ***, 2004).
- Ming, H. *et al.* Ultralow complexity long short-term memory network for fiber nonlinearity mitigation in coherent optical communication systems. [arXiv:2108.10212](https://arxiv.org/abs/2108.10212) (arXiv preprint) (2021).
- Kaneda, N. *et al.* Fpga implementation of deep neural network based equalizers for high-speed pon. In *Optical Fiber Communication Conference (OFC) 2020, T4D.2*. <https://doi.org/10.1364/OFC.2020.T4D.2> (Optical Society of America, 2020) (2020).
- Blalock, D., Ortiz, J. J. G., Frankle, J. & Gutttag, J. What is the state of neural network pruning? (2020). [arXiv:2003.03033](https://arxiv.org/abs/2003.03033).
- Han, S., Mao, H. & Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding (2016). [arXiv:1510.00149](https://arxiv.org/abs/1510.00149).
- Srinivas, S., Subramanya, A. & Babu, R. V. Training sparse neural networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* 455–462 (2017).

22. Hawks, B. *et al.* Ps and qs: Quantization-aware pruning for efficient low latency neural network inference. *Front. Artif. Intell.* <https://doi.org/10.3389/frai.2021.676564> (2021).
23. Sze, V., Chen, Y.-H., Yang, T.-J. & Emer, J. S. Efficient processing of deep neural networks: A tutorial and survey. *Proc. IEEE* **105**, 2295–2329. <https://doi.org/10.1109/JPROC.2017.2761740> (2017).
24. Liang, T., Glossner, J., Wang, L., Shi, S. & Zhang, X. Pruning and quantization for deep neural network acceleration: A survey. *Neurocomputing* **2101**, 09671 (2021).
25. Fujisawa, S. *et al.* Weight pruning techniques towards photonic implementation of nonlinear impairment compensation using neural networks. *J. Lightw. Technol.* <https://doi.org/10.1109/JLT.2021.3117609> (2021).
26. Li, M., Zhang, W., Chen, Q. & He, Z. High-throughput hardware deployment of pruned neural network based nonlinear equalization for 100-gbps short-reach optical interconnect. *Opt. Lett.* **46**, 4980–4983 (2021).
27. Oliari, V. *et al.* Revisiting efficient multi-step nonlinearity compensation with machine learning: An experimental demonstration. *J. Lightw. Technol.* **38**, 3114–3124 (2020).
28. Koike-Akino, T., Wang, Y., Kojima, K., Parsons, K. & Yoshida, T. Zero-multiplier sparse dnn equalization for fiber-optic qam systems with probabilistic amplitude shaping. In *2021 European Conference on Optical Communications (ECOC)*, 1–4 (IEEE) (2021).
29. Freire, P. J. *et al.* Transfer learning for neural networks-based equalizers in coherent optical systems. *J. Lightw. Technol.* **39**, 6733–6745. <https://doi.org/10.1109/JLT.2021.3108006> (2021).
30. Pelikan, M., Goldberg, D. E., Cantú-Paz, E. *et al.* Boa: The bayesian optimization algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference GECCO-99*, vol. 1, 525–532 (Citeseer) (1999).
31. Abadi, M. *et al.* TensorFlow: Large-scale machine learning on heterogeneous systems (2015). Software available from tensorflow.org.
32. Hoefer, T., Alistarh, D., Ben-Nun, T., Dryden, N. & Peste, A. Transfer learning for neural networks-based equalizers in coherent optical systems. *J. Mach. Learn. Res.* **2102**, 00554 (2021).
33. Allen-Zhu, Z., Li, Y. & Song, Z. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, 242–252 (PMLR) (2019).
34. Neill, J. O. An overview of neural network compression. [arXiv:2006.03669](https://arxiv.org/abs/2006.03669) (2020).
35. Neyshabur, B., Li, Z., Bhojanapalli, S., LeCun, Y. & Srebro, N. Towards understanding the role of over-parametrization in generalization of neural networks. [arXiv:1805.12076](https://arxiv.org/abs/1805.12076) (arXiv preprint) (2018).
36. Zhu, M. & Gupta, S. To prune, or not to prune: Exploring the efficacy of pruning for model compression. [arXiv:1710.01878](https://arxiv.org/abs/1710.01878) (arXiv preprint) (2017).
37. Hadidi, R. *et al.* Characterizing the deployment of deep neural networks on commercial edge devices. In *2019 IEEE International Symposium on Workload Characterization (IISWC)*, 35–48 (IEEE) (2019).
38. Yang, T.-J., Chen, Y.-H., Emer, J. & Sze, V. A method to estimate the energy consumption of deep neural networks. In *2017 51st Asilomar Conference on Signals, Systems, and Computers*, 1916–1920 (IEEE) (2017).
39. Närhi, M. *et al.* Machine learning analysis of extreme events in optical fibre modulation instability. *Nat. Commun.* **9**, 4923. <https://doi.org/10.1038/s41467-018-07355-y> (2018).
40. Agrawal, G. Chapter 2—pulse propagation in fibers. In *Nonlinear Fiber Optics (Fifth Edition)*, Optics and Photonics (ed. Agrawal, G.) 27–56 (Academic Press, Boston, 2013). <https://doi.org/10.1016/B978-0-12-397023-7.00002-4>.
41. Freire, P. J., Prilepsky, J. E., Osadchuk, Y., Turitsyn, S. K. & Aref, V. Neural networks based post-equalization in coherent optical systems: Regression versus classification. [arXiv:2109.13843](https://arxiv.org/abs/2109.13843) (arXiv preprint) (2021).
42. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (arXiv preprint) (2014).
43. Eriksson, T. A., Bülow, H. & Leven, A. Applying neural networks in optical communication systems: Possible pitfalls. *IEEE Photon. Technol. Lett.* **29**, 2091–2094 (2017).
44. Freire, P. J. *et al.* Neural networks-based equalizers for coherent optical transmission: Caveats and pitfalls. [arXiv:2109.14942](https://arxiv.org/abs/2109.14942) (arXiv preprint) (2021).
45. Matsumoto, M. & Nishimura, T. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.* **8**, 3–30 (1998).
46. Dong, X. & Zhou, L. Understanding over-parameterized deep networks by geometrization. [arXiv:1902.03793](https://arxiv.org/abs/1902.03793) (2019).
47. Bondarenko, A., Borisov, A. & Alekseeva, L. Neurons vs weights pruning in artificial neural networks. In *ENVIRONMENT. TECHNOLOGIES. RESOURCES. Proceedings of the International Scientific and Practical Conference*, vol. 3, 22–28 (2015).
48. Hu, H., Peng, R., Tai, Y. & Tang, C. Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. [arXiv:1607.03250](https://arxiv.org/abs/1607.03250) CoRR (2016).
49. Bartoldson, B., Morcos, A., Barbu, A. & Erlebacher, G. The generalization-stability tradeoff in neural network pruning. *Adv. Neural. Inf. Process. Syst.* **33**, 20852–20864 (2020).
50. Choukroun, Y., Kravchik, E., Yang, F. & Kisilev, P. Low-bit quantization of neural networks for efficient inference. [arXiv:1902.06822](https://arxiv.org/abs/1902.06822) (2019).
51. Yang, J. *et al.* Quantization networks. [arXiv:1911.09464](https://arxiv.org/abs/1911.09464) (2019).
52. Wu, H., Judd, P., Zhang, X., Isaev, M. & Micikevicius, P. Integer quantization for deep learning inference: Principles and empirical evaluation. [arXiv:2004.09602](https://arxiv.org/abs/2004.09602) (arXiv preprint) (2020).
53. Gholami, A. *et al.* A survey of quantization methods for efficient neural network inference. [arXiv:2103.13630](https://arxiv.org/abs/2103.13630) (arXiv preprint) (2021).
54. Hubara, I., Nahshan, Y., Hanani, Y., Banner, R. & Soudry, D. Accurate post training quantization with small calibration sets. In *International Conference on Machine Learning*, 4466–4475 (PMLR) (2021).
55. de Lima, T. F. *et al.* Machine learning with neuromorphic photonics. *J. Lightw. Technol.* **37**, 1515–1534 (2019).
56. Baskin, C. *et al.* Uniq: Uniform noise injection for non-uniform quantization of neural networks. *ACM Trans. Comput. Syst.* <https://doi.org/10.1145/3444943> (2021).
57. Freire, P. J. *et al.* Complex-valued neural network design for mitigation of signal distortions in optical links. *J. Lightw. Technol.* **39**, 1696–1705. <https://doi.org/10.1109/JLT.2020.3042414> (2021).
58. Albawi, S., Mohammed, T. A. & Al-Zawi, S. Understanding of a convolutional neural network. In *2017 International Conference on Engineering and Technology (ICET)*, 1–6 (Ieee) (2017).
59. Tran, N. *et al.* Ps and qs: Quantization-aware pruning for efficient low latency neural network inference. *Front. Artif. Intell.* **4**, 94 (2021).
60. Valladares, S., Toscano, M., Tufiño, R., Morillo, P. & Vallejo-Huanga, D. Performance evaluation of the nvidia jetson nano through a real-time machine learning application. In *International Conference on Intelligent Human Systems Integration*, 343–349 (Springer) (2021).
61. Tang, R., Wang, W., Tu, Z. & Lin, J. An experimental analysis of the power consumption of convolutional neural networks for keyword spotting. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5479–5483 (IEEE) (2018).
62. Holly, S., Wendt, A. & Lechner, M. Profiling energy consumption of deep neural networks on nvidia jetson nano. In *2020 11th International Green and Sustainable Computing Workshops (IGSC)*, 1–6 (IEEE) (2020).
63. Kaup, F., Gottschling, P. & Hausheer, D. Powerpi: Measuring and modeling the power consumption of the raspberry pi. In *39th Annual IEEE Conference on Local Computer Networks*, 236–243 (IEEE) (2014).

Acknowledgements

SKT and MKK are partially supported by the EPSRC programme Grant TRANSNET, EP/R035342/1. PJF and DAR acknowledge the support from the EU Horizon 2020 Marie Skłodowska-Curie Action projects No. 813144 (REAL-NET) and 860360 (POST-DIGITAL), respectively. JEP and SKT acknowledge the support of the Leverhulme Trust project RPG-2018-063.

Author contributions

D.A.R., P.J.F., and J.E.P. conceived the study. D.A.R. and P.J.F. proposed the neural network model. D.A.R. performed the numerical simulations, designed the experimental set-up and obtained the experimental results. P.J.F. generated the data and performed the architecture optimization. D.A.R. and P.J.F. designed the figures and tables. D.A.R., P.J.F., and J.E.P. wrote the manuscript, with the assistance of M.K.K. and S.K.T. All authors reviewed the manuscript. The work of D.A.R. was supervised by M.K.K. and S.K.T. The work of P.J.F. was supervised by J.E.P., A.N. and S.K.T.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to D.A.R. or S.K.T.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022