

Research Article

Ilker Oguz*, Louis J.E. Suter, Jih-Liang Hsieh, Mustafa Yildirim, Niyazi Ulas Dinc, Christophe Moser and Demetri Psaltis

Training hybrid neural networks with multimode optical nonlinearities using digital twins

<https://doi.org/10.1515/nanoph-2025-0002>

Received January 2, 2025; accepted March 18, 2025;

published online April 2, 2025

Abstract: The ability to train ever-larger neural networks brings artificial intelligence to the forefront of scientific and technical discoveries. However, their exponentially increasing size creates a proportionally greater demand for energy and computational hardware. Incorporating complex physical events in networks as fixed, efficient computation modules can address this demand by decreasing the complexity of trainable layers. Here, we utilize ultrashort pulse propagation in multimode fibers, which perform large-scale nonlinear transformations, for this purpose. Training the hybrid architecture is achieved through a neural model that differentially approximates the optical system. The training algorithm updates the neural simulator and back-propagates the error signal over this proxy to optimize layers preceding the optical one. Our experimental results achieve state-of-the-art image classification accuracies and simulation fidelity. Moreover, the framework demonstrates exceptional resistance to experimental drifts. By integrating low-energy physical systems into neural networks, this approach enables scalable, energy-efficient AI models with significantly reduced computational demands.

Keywords: optical computing; digital twins; neural networks; artificial intelligence; fiber optics; optimization algorithms

1 Introduction

Modern artificial intelligence (AI) models are currently driving a transformation across a wide variety of fields including science, industry, healthcare, and education with their unprecedented ability to learn an immense scale of information and solve complex problems. These capabilities have been mainly enabled by two key factors: algorithmic breakthroughs in training neural networks with increasingly larger numbers of parameters and datasets, and the advent of massively parallelized digital hardware to execute these algorithms. So far, employing larger neural networks constantly provided higher AI accuracies in various tasks such as language or image processing, given correspondingly large training sets [1]. The observation of this trend has also led to state-of-the-models with exponentially larger number of parameters, nearly doubling every year [2]. However, this rapid increase in model sizes, coupled with the broader adoption of AI, has created a compounding effect, doubling this technology's overall energy consumption every 100 days [3] and resulting in a correspondingly larger environmental impact due to carbon emissions [4].

A promising approach to performing complex tasks without increasing the number of trainable parameters is the addition of random features, where a subset of weights is randomly assigned instead of trained [5], [6]. The effectiveness of this method has been demonstrated within architectures such as deep convolutional networks [7], transformers [8], [9], and graph neural networks [10]. Despite their potential, random features are not universally integrated into modern neural networks, because they still require significant computation contributing to energy consumption and latency on conventional hardware. To address this challenge, low-power, high-speed, and high-dimensional physical systems offer a promising alternative for efficiently generating large numbers of fixed weights [11]–[14]. While fixed physical layers can be used as initial preprocessing in a straightforward manner by training the following digital layers conventionally, adding programmable layers before

Ilker Oguz and Louis J.E. Suter contributed equally to this work.

*Corresponding author: **Ilker Oguz**, EPFL, Institute of Electrical and Micro Engineering, 1015 Lausanne, Switzerland,

E-mail: ilker.oguz@outlook.com. <https://orcid.org/0000-0003-2674-684X>

Louis J.E. Suter, Jih-Liang Hsieh, Mustafa Yildirim, Niyazi Ulas Dinc, Christophe Moser and Demetri Psaltis, EPFL, Institute of Electrical and Micro Engineering, 1015 Lausanne, Switzerland.

<https://orcid.org/0000-0002-9606-8098> (N.U. Dinc).

<https://orcid.org/0000-0002-2078-0273> (C. Moser)

the physical layer is more challenging, especially for highly complex, nonlinear systems. One such complex system is the propagation of pulses in MMFs, where tight confinement of light over long distances can create nonlinear modal couplings with low power in addition to the linear transformation due to propagation [15], [16]. Backpropagation-free methods, such as genetic algorithms [17], surrogate optimization [18], and local learning [19], have shown promise in significantly reducing parameter counts while achieving competitive accuracy by treating optical systems as fixed black boxes. Nevertheless, further research is needed to scale these approaches for more demanding applications requiring large number of parameters.

On the other hand, to be able to use the error backpropagation algorithm, which successfully trains today's AI models reaching trillions of weights, accurate knowledge of each layer's analytical representation is required, yet this is challenging to achieve with complex physical layers, as shown in Figure 1(a). Instead, approximating their functionality with neural networks trained with examples from the physical system allowed different applications with intricate optical systems including telecommunication [20], femtosecond laser inscription [21], image acquisition [22], and projection [23] alongside training physical AI models [13].

Even if a model can predict the behavior of a subset of all possible samples, it is highly challenging to fully

generalize efficient approximators to high-dimensional nonlinear systems such as pulse propagation in multimode fibers. Moreover as in many analog information processing systems [24], [25] drift and instabilities in the compute characteristics is eventually unavoidable in MMF-based systems as well, while methods such as temperature and mechanical stabilization [26] and wavefront shaping [27] are shown to alleviate these issues.

Here, we demonstrate an approach that can introduce nonlinear optical layers to neural networks in a scalable manner and can train them effectively while being resistant to the deteriorating effects mentioned. To achieve this goal, we synergistically update a digital twin of the optical system throughout the training process and approximate the Jacobian matrix of the system, which holds the partial derivatives between its output and input channels. This way, the layers preceding the physical ones can receive feedback on the effect of their weights on the output and can be updated, as indicated in Figure 1(b). Approximation of the gradients over the physical layers allows for the utilization of them in multiple locations, creating large-scale networks with only a small portion of the weights trainable and digital as in Figure 1(c), while the forward process can fully benefit from the optical weights without any additional digital operations during the inference phase. Our experimental results, obtained with a combination of digital and physical layers in the 3-layer system shown in Figure 2(a), quantitatively

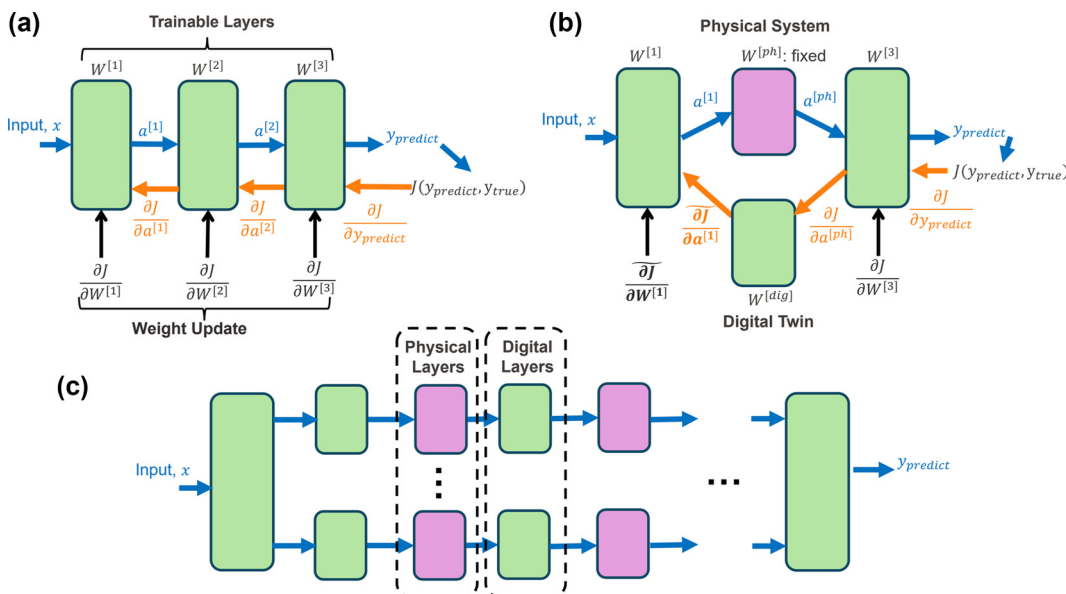


Figure 1: The information flow during conventional training of neural networks, and training scheme with physical layers. (a) In fully digital neural networks, an analytical model of the functionality of each layer is precisely known and their derivatives are used in the calculation of the updates to the weights to minimize error. (b) Physical layers can be included in the forward computation in neural networks while their digital twin serves as a differentiable function to replace them in the backpropagation step. The data acquired in the forward pass serve as training dataset that constantly refines the digital twin. (c) Gradient approximation with the digital twin allows for large scale hybrid networks containing physical layers in different locations.

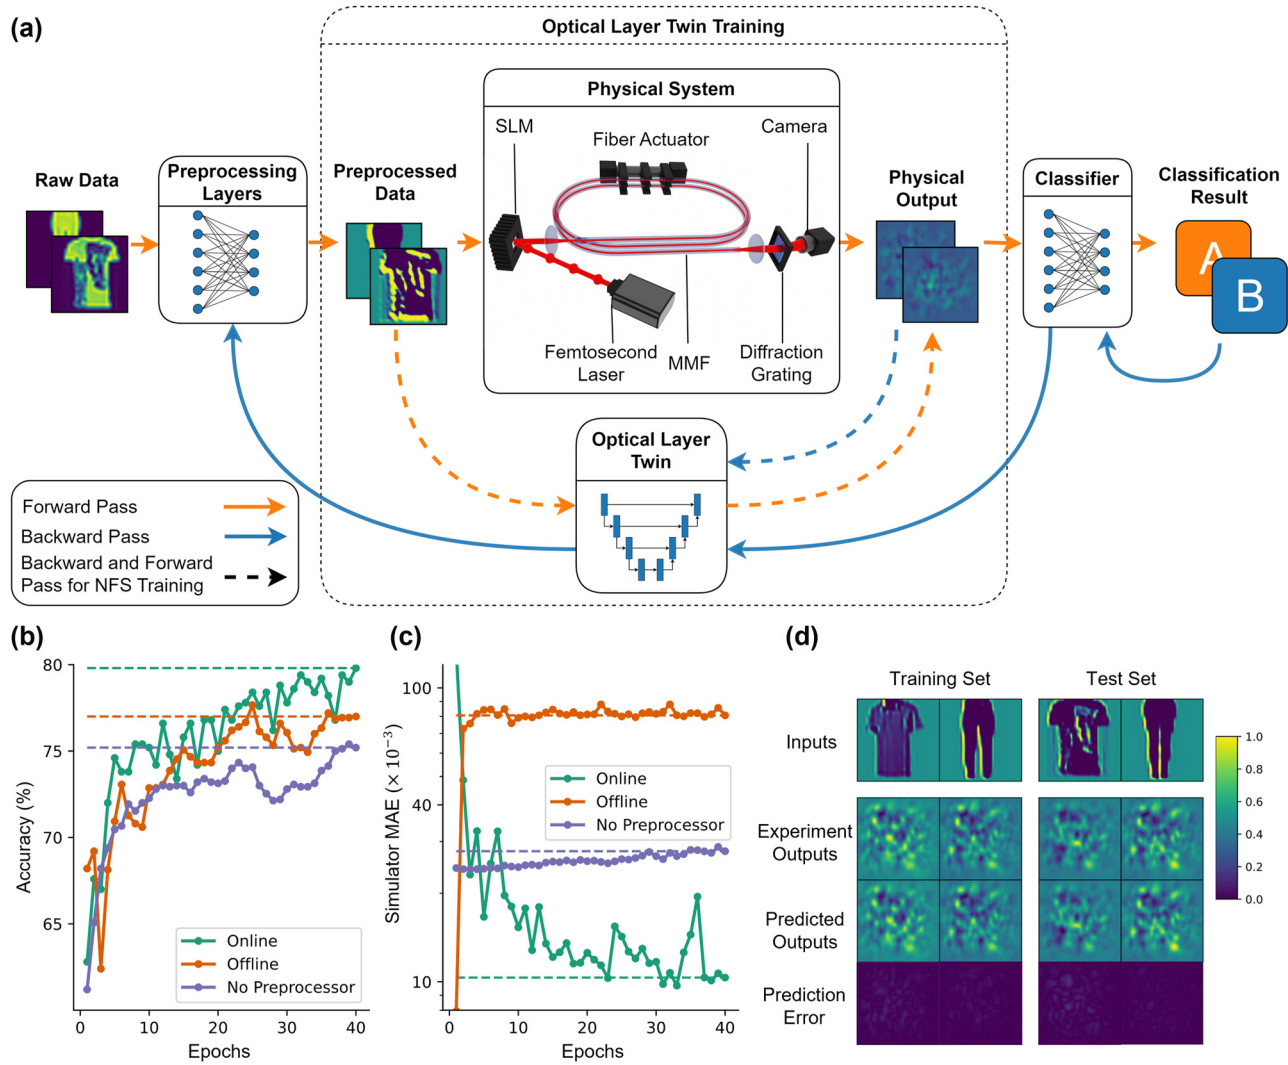


Figure 2: The architecture, algorithm, and outcomes of training a neural network with a physical layer. (a) The data flow in the network through digital and physical layers both in forward and error backpropagation steps. The physical layer consists of an optical system based on the propagation of spatially modulated ultrashort pulses in a multimode fiber. (b) The progress of test accuracy over training epochs for different learning structures on Fashion-MNIST dataset. (c) The mean absolute error (MAE) between Optical Layer Twin (OLT) predictions and experimentally obtained intensity patterns for image in the test set for different training procedures. (d) Some intensity distribution samples from OLT, experimental system, and their absolute differences for the same input images.

demonstrate the effectiveness of the method for AI tasks in terms of accuracy, resistance to drifts, and simulation fidelity.

2 Methods and results

2.1 Guiding model training with the Optical Layer Twin

Our approach leverages a digital twin, named Optical Layer Twin (OLT), to approximate the physical response of propagation through the MMF, a high-dimensional nonlinear

system. The propagation of spatially phase-modulated ultrashort pulses (10 ps long at 1,030 nm with 125 kHz repetition rate, 12.6 mW fiber-coupled average power, generated by Amplitude Laser, Satsuma) in a multimode fiber (OFS, bend-insensitive OM2, 50 μm core diameter, 0.20 NA, 5 m long) constitutes the physical event in the optical system. As shown in Figure 2(a), the input information enters the system as a 2-dimensional phase distribution on the spatial light modulator (SLM) (Meadowlark HSP1920), which determines the modes coupled into the fiber. The propagation of the modulated beam through the silica core of the fiber induces spatiotemporal nonlinearities due to $\chi^{(3)}$

light–matter interactions such as cross-phase modulation and four wave mixing, remapping the initial field distribution to a completely different pattern, which is measured finally on a camera. Optical power coupled to the MMF is selected to be the optimal level for obtaining rich nonlinearities without impairing the performance with beam clean-up effects as previously established in the literature [18]. Appendix Note 1 explains further these nonlinear propagation dynamics. The fiber actuator can change the physical conformation of the fiber controllably and can serve as a source of reproducible perturbation of the physical system. Before reaching the camera, a diffraction grating disperses the beam spatially separating its wavelength components. This enriches the transformation between the input phase and the output intensity patterns by distributing the spectral effects to camera pixels.

The physical system provides a large set of nonlinear connections to the overall neural network model, which can learn to tackle various machine learning problems through trainable initial and final layers. In the demonstration shown in Figure 2, the initial preprocessing layer is a single convolutional layer, while the final classifier is a fully connected layer. During training, OLT accompanies the forward computation path of the model, including the physical system, by approximating the physical response with a multilayer neural network, which is a differentiable function for the error signal's backward propagation. OLT itself is a standalone neural network that predicts the output intensity distribution of the physical system given the phase modulation pattern on the SLM. Since incorporating the OLT into the training loop enables updating the weights of the layers preceding the physical model with gradients from the downstream task, it does not add any extra overhead to the forward information flow in the model. Consequently, the computational complexity of inference operations after training remains unchanged. Furthermore, the OLT is continuously updated with experimental data generated during operation to achieve the highest approximation fidelity. We refer to this mode of operation as “Online Learning” in this paper and compare it with an approach where the OLT is preconfigured and kept fixed throughout training, termed “Offline Learning.” In Figure 2(a), the operations shown with solid lines take place in both approaches, whereas online learning also applies update steps indicated with dashed lines. The average gradient of the loss across all samples in a batch determines the weight update direction and magnitude. Due to the limited memory space on the GPU, the batch size is selected to be 15. This value also provides a balanced trade-off between generalization performance

and computation time. The details of these training algorithms are provided in Appendix Note 2.

In scenarios where the preprocessing layers are more complex and its weights significantly alter the input data sent to the physical system, online learning becomes increasingly crucial. This is because a fixed OLT, trained on a limited or specific dataset, may fail to generalize to new input distributions introduced by updated preprocessing weights. The wide gap between the dataset used to train the OLT and the new, altered input samples it processes later can reduce the fidelity of the OLT, leading to imprecise updates to the upstream layers in training.

We compared computational performances of the two learning algorithms on a subset of Fashion-MNIST dataset [28] with only 1,500 training samples, also including a simpler approach that processes raw data directly with the physical system and classifies its outputs without any preprocessing. As shown in Figure 2(b), adding a preprocessing layer and training with offline learning improves the classification accuracy on 500 test samples from 75 % to 77 %. When the training algorithm keeps updating the OLT throughout training, a better performing preprocessor brings the accuracy to 80 %. Thus, the task accuracy improves 5 % compared to the case without any preprocessing, which can be considered a photonic extreme learning machine, a fixed and large-scale nonlinear transform following a single linear trainable layer [29]. Furthermore, the presented approach yields approximately the same performance with feedforward diffractive networks, while being trained on a significantly smaller portion of the dataset, and using about 16 thousand weights compared to 400 thousand in the diffractive case [30].

The fidelity of OLT to the actual experiment, visualized in Figure 2(c), explains the obtained accuracy levels. Without a preprocessing layer, MAE increases slightly, about 10 % over the course of 40 epochs taking approximately 2 h. This increase stems directly from slow drifts in the experimental system over time, since the raw experimental output data and the predictions of the pretrained and fixed OLT are utilized to calculate it. In the offline learning experiment, the error level rises rapidly during the initial epochs, as changes in the preprocessing layer alter the distribution of inputs to the physical system, which the fixed OLT cannot accommodate, and MAE reaches nearly 10^{-1} . Here, epoch means one pass over the entire dataset during training. In contrast, with online learning, the OLT tracks these distributional changes, ultimately achieving high precision in predicting experimental outputs. The precision level reaches a normalized MAE of 1.03×10^{-2} , meaning that signal-to-noise (SNR) ratio of the OLT predictions is 96.8. The examples in

Figure 2(d), taken from the final epoch of the online training procedure, demonstrate the effectiveness of the OLT, with nearly visually identical speckle patterns between the experimental outputs and the blind predictions of the neural simulator. Remarkably, the OLT achieves these accurate predictions with a latency of 30 ms, while an analytical simulation of the wave propagation dynamics, using the multimode nonlinear Schrödinger equation for the same fiber, would take approximately 500 s on a GPU, even when limited to only 15 of the fiber's 240 modes [31].

2.2 Effects of OLT architecture on modeling fidelity and computational costs

Capturing the complex dynamics of multimode nonlinear fiber optics with a neural network requires a large-scale, deep architecture. In this study, the mapping to be learned is from the two-dimensional input phase distribution on the SLM to the two-dimensional output intensity on the camera. A convolutional U-Net [32] architecture is particularly well-suited for this image-to-image translation task. It employs

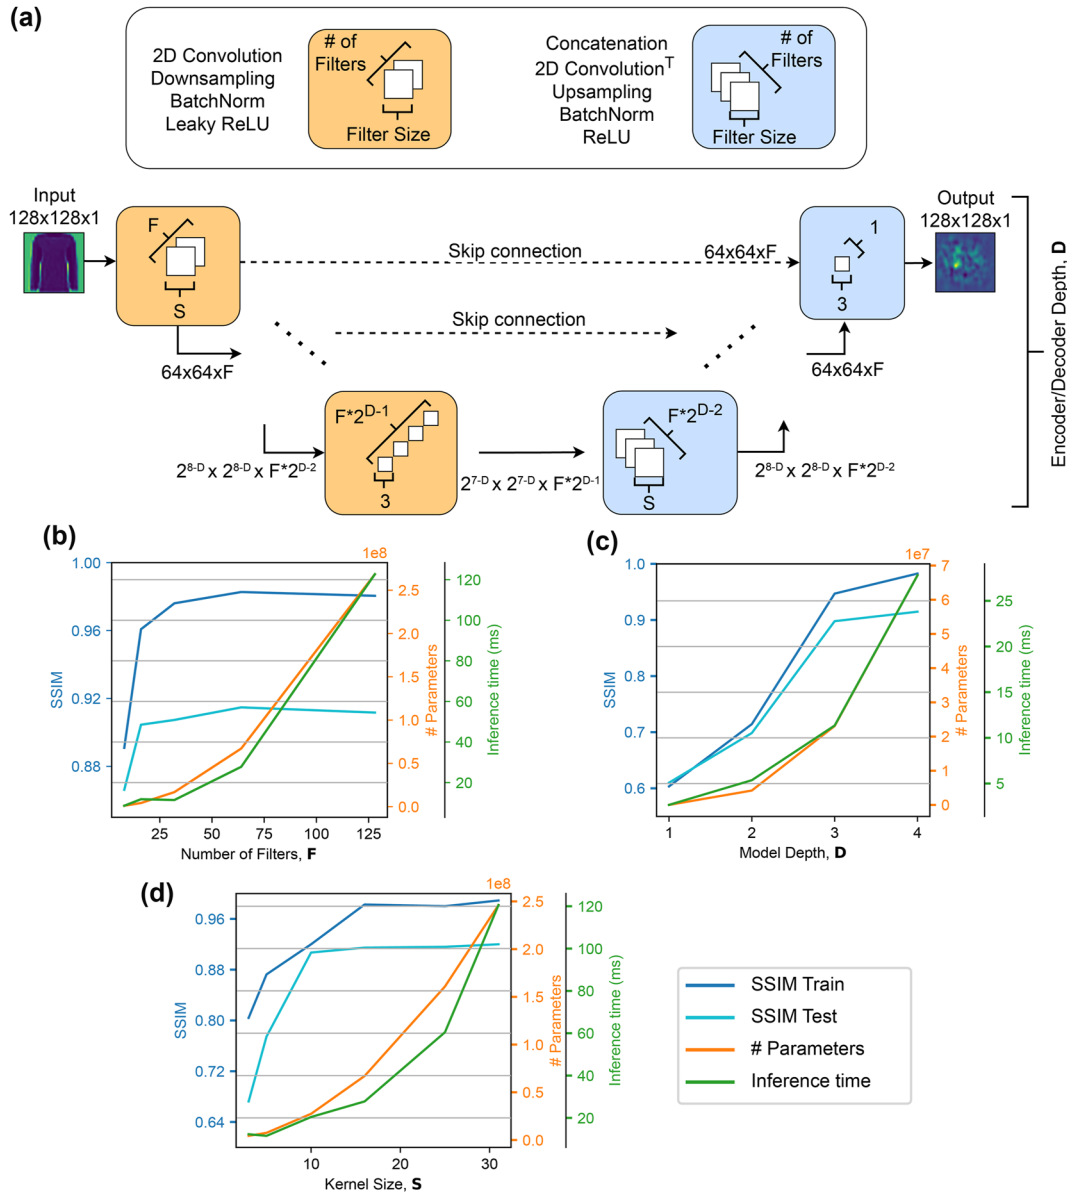


Figure 3: Architecture of the OLT and its scaling behavior under various hyperparameter configurations. (a) Network structure of the OLT, based on the U-Net architecture. The layers forming the downscaling (orange) and upscaling (blue) blocks, along with their dimensionalities, are shown. (b) Impact of the number of convolutional filters on structural similarity index (SSIM), total model parameters, and inference time per sample. (c) Effect of varying the number of downscaling and upscaling layers, which defines the model depth, on the same metrics. (d) Influence of convolutional kernel sizes on SSIM, model parameters, and inference time.

encoding blocks to downscale and extract features into a latent space, and decoding blocks to reconstruct the desired output from these latent features through upscaling, as visualized in Figure 3(a).

Optimizing the scale of this architecture for the task at hand requires careful analysis, as a smaller or disproportionate network might not be able to fully capture the underlying dynamics, while an overly complicated structure might tend to overfit or require excessive amount of compute resources. Using the dataset of inputs and outputs to the optical system, we trained different OLT architectures by varying the depth, defined by the number of encoding and decoding blocks and denoted as “D” in Figure 3. The improvement in the structural similarity index (SSIM) between OLT predictions and optical measurements, as shown in Figure 3(c), highlights the clear benefit of deeper

architectures. A higher SSIM means a better prediction fidelity of experimental results, corresponding to lower MAE and resulting in a more precise gradient estimation by the OLT.

However, for architectures with 4 down- and up-sampling blocks, increasing the number of convolutional filters per block, denoted as “F,” above 16 creates significantly larger models and, more critically, longer inference times, without improving the accuracy, as demonstrated in Figure 3(b). Similarly, while increasing the size of convolutional kernels, shown as “S,” up to 16×16 improves accuracy, larger kernel sizes do not yield further gains, despite resulting in a much larger architecture as plotted in Figure 3(d). Under the light of these findings, we identified the ideal OLT architecture for the given optical system as a U-Net that accepts and generates 128×128 pixels images with

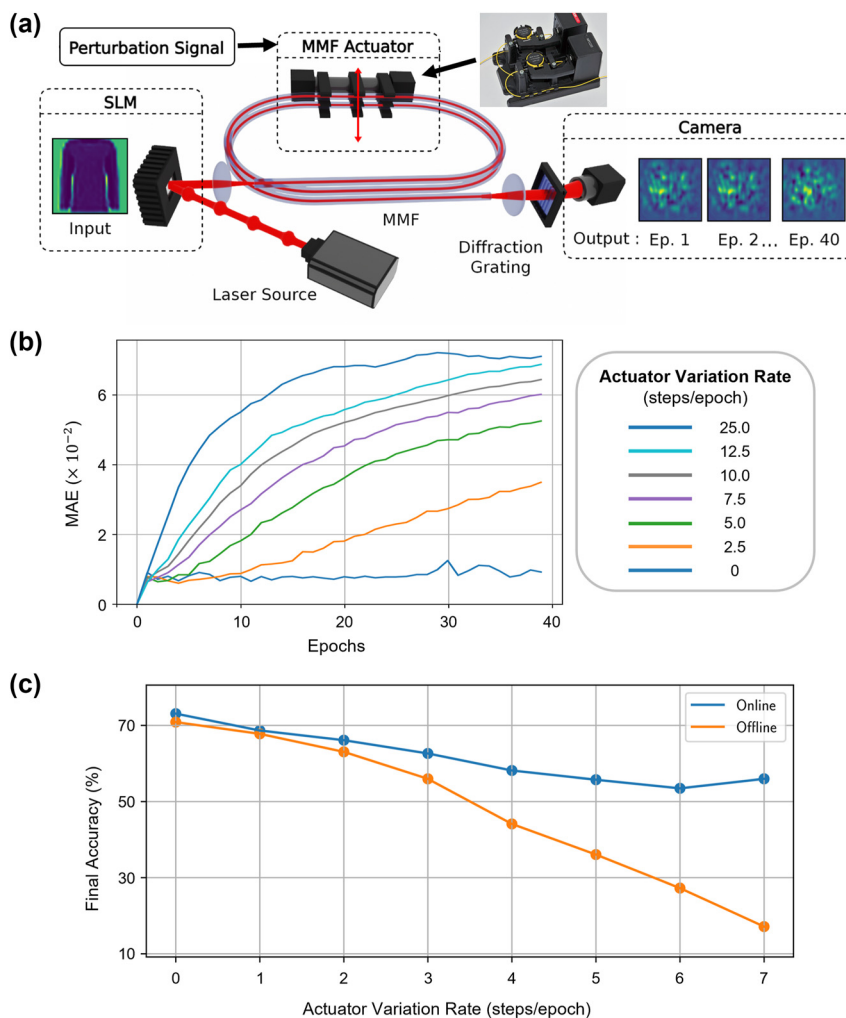


Figure 4: The method for applying dynamic perturbations to a multimode fiber based optical system, their effects on the output statistics, and accuracy of the overall network. (a) A mechanical actuator is added to bend the MMF in the setup shown in Figure 2(a), introducing dynamic perturbations. The effects are observed through changes in output speckles at each epoch for the same input pattern. (b) Comparison of the dataset's outputs in the first epoch versus subsequent epochs, quantified using the average MAE at different actuator variation rates. (c) Dependency of final test accuracies on the perturbation rate, shown for offline and online learning schemes.

4 encoding and decoding blocks, starting with 16 convolutional kernels of 16×16 parameters.

2.3 Learning and computing with a physical system under dynamic perturbation

As in various analog information processing systems, gradual small changes in the working characteristics of optical physical layers can accumulate and significantly impact final predictions. To investigate these effects and their impact on the neural network performance in an accelerated manner, a mechanical actuator is incorporated into the optical system shown in Figure 2(a). The mechanical actuator (Thorlabs MPC320), originally designed to create stress-induced birefringence in single-mode fibers, rotates a small portion of the MMF with respect to the rest, inducing controllable perturbations to linear mode couplings. The changes in output speckle shapes observed in Figure 4(a) result from these differences in mode couplings.

Rotating the actuator by a set number of steps, each corresponding to 0.12° , at the beginning of every epoch allows quantitative analysis of its effect on the dataset's representation at the optical system's output, as shown in Figure 4(b). This analysis demonstrates that, without actuator movement, the system's characteristics remain stable in the experiment's timeframe. In contrast, induced perturbations increase the differences in speckle distributions linearly at smaller rates and before saturating at higher levels.

When these varying rates of externally induced perturbation are applied to the physical layer during the training, the online training method provides larger final accuracy advantages with faster drifts. Results in Figure 4(c) indicate this advantage to reach 39 % final accuracy improvement. Also, even without any active rotation, this experiment yields lower task accuracies compared to the experiment without the actuator in Figure 2(a). This reduction is due to high-order mode filtering caused by the tight winding of the MMF on the actuator, which has a bending diameter of 18 mm. This well-known effect induces higher losses for higher-order modes and couples them out of the fiber, observed with bending diameters starting from 40 mm [33]. The reduced task accuracy with fewer effective modes in the fiber underscores the importance of high dimensionality in physical layers.

3 Conclusions

In this study, we studied an MMF-based optical system as a prototype for integrating complex, nonlinear physical

phenomena as fixed layers into neural networks. To enable the training of such networks using the error backpropagation algorithm, which facilitated the artificial intelligence models to reach their current immense scales, a method for estimating the derivatives of the physical layer's outputs is required. This becomes crucial when the underlying physical relations are intractably complex for analytical computation.

Our approach employs an image-to-image translating neural network to model the dynamics of the optical system. This differentiable neural network approximates the system's outputs and also enables the calculation of gradients required for backpropagation. Through extensive exploration of hyperparameter configurations, we identify that a neural network with 70 million parameters, achieving an inference time of 30 ms on an NVIDIA RTX 3070 Ti GPU, provides the optimal trade-off between precision and computational cost, attaining an SNR of approximately 100. Furthermore, by continuously updating this proxy model, the method ensures accurate gradient predictions throughout the training procedure. Continuous updates to the model are particularly necessary when the distribution of inputs to the physical system shifts due to the weight updates in preceding layers. Figure 5 illustrates the evolution of the optical system's and the OLT outputs over the training process. Notably, if the OLT is trained exclusively on raw sample inputs, changes in the input distribution can lead to reduced prediction accuracy. However, online learning effectively mitigates this issue, maintaining accuracy under evolving input distributions. While training with random inputs and their optical outputs can enhance the OLT's coverage, it reduces precision in specific subspaces due to the model's finite capacity.

Environmental and device-related factors can induce subtle changes in the transformation characteristics of

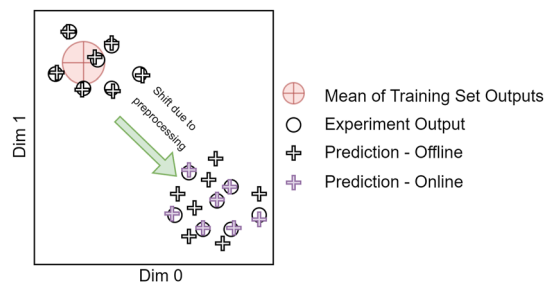


Figure 5: The visualization of the changes in the optical system's and OLT outputs during the training process due to the updates to the preprocessing layers.

analog systems. To investigate these effects, we artificially accelerated the system's drifts by incorporating a mechanical stage into the optical system, simulating systematic perturbations. Our findings reveal that while increased perturbations degrade neural network classification accuracy, updating the neural surrogate significantly mitigates this impact, improving task accuracy by nearly 40 %.

This study demonstrates the symbiosis of a very large-scale optical and smaller scale digital layers in a neural network. Since the employed training algorithm, error back-propagation, has a strong track record in scaling up learning systems, the proposed approach paves the way for competitive, physics-assisted AI models with minimal digital computation needs in inference, which generally consumes the vast majority of energy and time within the lifetime of a conventional model. Depending on the variation rate of the utilized physical mechanism, the weights can be updated once in many inference operations, such that the overhead due to the OLT is negligible and the overall model keeps optimal performance constantly. Moreover, for more complex tasks such as language or image generation, deeper architectures with multiple physical layers stacked in width and depth can be constructed using the principles outlined here.

We utilized optoelectronic devices to bridge physical layers with the digital domain and incorporate into neural networks, although physical interactions in optical data links, such as those in MMFs, are readily available and often treated as errors to be corrected. The proposed approach can reposition them as computational units by modeling these interactions with the OLT and incorporating them into neural networks. This enables existing optical data to be leveraged for computation, providing a pathway to fewer digital parameters and enhance AI efficiency, offering energy-efficient, low-latency workflows for environments like data centers.

Research funding: This study was supported by the Swiss National Science Foundation, Sinergia grant CRSII5_216600 LION: Large Intelligent Optical Networks.

Author contributions: IO, J-LH and MY prepared the experimental system. LS and NUD developed the codebase. IO and LS compiled the data. CM and DP supervised the project. All authors contributed to the preparation of this manuscript, have accepted responsibility for its entire content, and approved its submission.

Conflict of interest: Authors state no conflicts of interest.

Data availability: Data used and produced in this study is available from the corresponding author upon reasonable request.

Appendix

Appendix Note 1: Nonlinear multimode fiber propagation

For an ideal fiber with a perfectly homogenous distribution of refractive index, light coupled to any mode propagates throughout on the same mode only with a phase change. However, perturbations to the ideal refractive index distribution and nonlinear propagation can cause the coupling of light from one mode to the other. This coupling can be formulized by using perturbation theory. The coupling coefficient is related to the overlap integral of the permittivity perturbation, $\Delta\epsilon(x, y, z)$, and the coupled mode shapes, $E_n(x, y)$, with a constant coefficient of p :

$$C_{p,n}(z) = p \int \Delta\epsilon(x, y, z) E_p^*(x, y) E_n(x, y) dx dy$$

In addition to the linear intermodal coupling due to waveguide perturbation, which simply can be induced by the bending of the fiber, nonlinear effects due to the light-matter interaction inside the waveguide can be represented as a mode-coupling mechanism. In general, these nonlinearities are represented as higher-order susceptibilities of the material and can be grouped as either parametric or nonparametric [34]. Parametric nonlinear processes do not change the quantum state of the material; hence, they do not exchange energy with the material, and compared to nonparametric processes, they occur instantly. One of the most common optical parametric processes is the Kerr effect, manifesting itself in events such as self-phase modulation and cross-phase modulation, creating a phase shift and spectrum change [35]. On the other hand, nonparametric nonlinear optical processes create another set of effects on the propagation of intense light such as stimulated Raman and Brillouin scattering. In these events, the interaction of photons with matter creates and destroys phonons, changing photons' energy and momentum. Thus, through these processes, optical beams of different wavelengths are generated and amplified.

These linear and nonlinear scattering events can be concisely formulated as coupling between modes of a multimode fiber. The following equation illustrates the evolution of each mode coefficient as it propagates through a

multimode fiber by accounting for the effects of dispersion, linear coupling due to bending and other refractive index

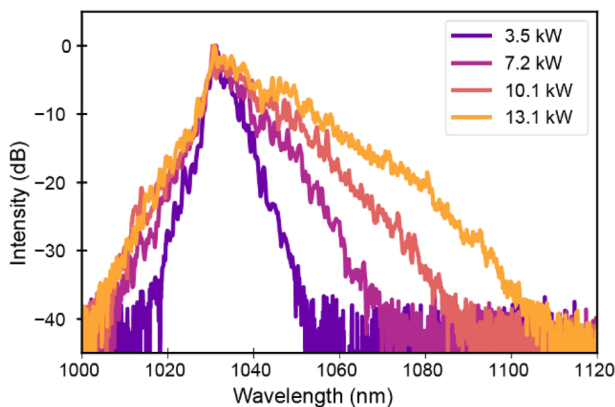
perturbations, and nonlinear coupling due to the Kerr effect and Raman scattering [36].

$$\begin{aligned} \frac{\partial A_p}{\partial z} = & \underbrace{i\delta\beta_0^p A_p - \delta\beta_1^p \frac{\partial A_p}{\partial t} - i\frac{\beta_2}{2} \frac{\partial^2 A_p}{\partial t^2}}_{\text{Dispersion}} + \underbrace{i \sum_n C_{p,n} A_n}_{\text{Linear Mode Coupling}} \\ & + \underbrace{i\frac{\gamma}{3} \left(1 + \frac{i}{\omega_0} \frac{\partial}{\partial t}\right) \sum_{l,m,n} \eta_{plmn} \times \left[\underbrace{(1-f_R)A_l A_m A_n^*}_{\text{Kerr contribution}} + \underbrace{f_R A_l \int h_R(\tau) A_m(z, t-\tau) A_n^*(z, t-\tau) d\tau}_{\text{Raman contribution}} \right]}_{\text{Nonlinear Mode Coupling}} \end{aligned}$$

Here, γ is the nonlinear coefficient, f_R is the Raman effect ratio in the nonlinear effect, $h_R(\tau)$ is the impulse response of the Raman scattering process, and η_{plmn} is the nonlinear mode coupling tensor, containing the amount of interaction for a set of 4 modes, and defined as:

$$\eta_{p,l,m,n} = \frac{\iint dx dy E_p E_l E_m E_n}{[\iint dx dy E_p \iint dx dy E_l \iint dx dy E_m \iint dx dy E_n]^2}.$$

In our experiments, we observed the ideal computing capabilities to occur at the level where maximum Kerr nonlinearities happen while Raman scattering is still negligible, which occurs around 10 kW of pulse peak power, as shown in Appendix Figure 1 [18]. For silica core, the first Raman peak is expected to be around 1,080 nm for the pump around 1,030 nm. The average fiber coupled power of 12.6 mW in this study is selected to satisfy this condition.



Appendix Figure 1: Output spectra of the multimode fiber under different pulse peak power levels, reused from [18].

Appendix Note 2: Online training algorithm

In this methodology, the forward pass involves processing input x through a series of layers and functions: the pre-processing layer (f_p), the optical system (f_{OS}), the classifier (f_c), and the loss function (f_l). The overall operation can be expressed as: $(f_l \circ f_c \circ f_{OS} \circ f_p)(x)$. During the backward pass, the optical system (f_{OS}) is replaced by its differentiable proxy, the Optical Layer Twin (f_{OLT}), enabling gradients to flow back through the system.

$$\begin{aligned} \nabla(f_l \circ f_c \circ f_{OS} \circ f_p)(x) &\approx J_{f_p}(x)^T J_{f_{OLT}}(f_p(x))^T \\ &\times J_{f_c}(f_{OS}(f_p(x)))^T \nabla f_l(f_c(f_{OS}(f_p(x)))). \end{aligned}$$

After gradient of the loss function is calculated, trainable layers are updated accordingly, for instance for the preprocessor, the update amount, $\Delta\theta_p$, is:

$$\Delta\theta_p \approx -\eta \nabla_{\theta_p}(f_l \circ f_c \circ f_{OS} \circ f_p)(x).$$

As it is also evident in this notation, the OLT's main functionality is to allow approximation of the Jacobian of the optical system such that $J_{f_{OLT}} \approx J_{f_{OS}}$. TensorFlow's `@custom_gradient` decorator allows the forward and backward computations to follow separate paths by using a single function, with the gradient of the experiment's output being explicitly defined. A key consideration is the use of vector-Jacobian products (VJP) instead of full Jacobian matrices, as the latter would be computationally prohibitive for high-dimensional data. For example, if the experiment maps a batch of 10 images of size 128×128 to equivalent outputs, the Jacobian will contain $(10 \cdot 128 \cdot 128)^2 \approx 1.7 \times 10^{10}$ elements. With 32 bit floating-point precision, this would require over 100 GB of memory, making it impractical for training. Instead, the VJP computes the gradient efficiently by combining the upstream gradient vector u with the output of the OLT:

$$\begin{aligned}\text{custom_gradient}(u) &= \text{vjp}(u, f_{OLT}) = \nabla u \cdot f_{OLT} \\ &= u^T J_{f_{OLT}}.\end{aligned}$$

At every step of training, the weights of the layers preceding and succeeding the optical layer (preprocessing and classifier layers) are updated following this formalism. In the “Online Learning” approach, each step also refines the OLT by using the data obtained in the forward pass from the experiment. For the same inputs, let’s call the output of the optical system as y_{OS} and the output of the OLT as y_{OLT} , then the loss function for the refinement step can be defined as the squared error: $L_{\text{refine}} = |y_{OS} - y_{OLT}|^2$. The derivative of the loss function with respect to the parameters of the OLT provides the update direction and amount: $\Delta\theta_{OLT} = -\eta \frac{\partial L_{\text{refine}}}{\partial \theta_{OLT}}$, η being the learning rate. This refinement step closely matches the behavior of the OLT to the experimental system, as the input distribution evolves due to updates in the preprocessing layer and possible physical drifts. This ensures high fidelity weight updates, reducing discrepancies and stabilizing the training process.

The presented experimental results are obtained with a preprocessing block of 6 linear convolutional layers of 1 kernel with 6×6 parameters, followed with a single sigmoid nonlinearity, and a classifying layer flattening 40×40 images and predicting the output with 10 softmax neurons. A stochastic gradient descent algorithm with 10^{-3} learning rate optimized the trainable parameters.

References

- [1] T. Henighan, *et al.*, “Scaling laws for autoregressive generative modeling,” *arXiv*: arXiv:2010.14701, 2020. <https://doi.org/10.48550/arXiv.2010.14701>.
- [2] Data Page: *Parameters in notable artificial intelligence systems*, part of the following publication: Charlie Giattino, Edouard Mathieu, Veronika Samborska and Max Roser (2023) - “Artificial Intelligence”. Data adapted from Epoch. [Online]. Available at: <https://ourworldindata.org/grapher/artificial-intelligence-parameter-count> [accessed: Dec. 09, 2024].
- [3] S. Zhu, *et al.*, “Intelligent computing: the latest advances, challenges, and future,” *Intell. Comput.*, vol. 2, p. 6, 2023.
- [4] A. S. Luccioni, Y. Jernite, and E. Strubell, “Power hungry processing: watts driving the cost of AI deployment?” in *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 2024, pp. 85–99.
- [5] S. Scardapane and D. Wang, “Randomness in neural networks: an overview,” *WIREs Data Min. Knowl. Discov.*, vol. 7, no. 2, p. e1200, 2017.
- [6] F. Liu, X. Huang, Y. Chen, and J. A. K. Suykens, “Random features for kernel approximation: a survey on algorithms, theory, and beyond,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 10, pp. 7128–7148, 2022.
- [7] J. Frankle, D. J. Schwab, and A. S. Morcos, “Training BatchNorm and only BatchNorm: on the expressive power of random features in CNNs,” *arXiv*: arXiv:2003.00152, 2021. <https://doi.org/10.48550/arXiv.2003.00152>.
- [8] K. M. Choromanski, *et al.*, “Rethinking attention with performers,” in *International Conference on Learning Representations*, 2020.
- [9] H. Peng, N. Pappas, D. Yogatama, R. Schwartz, N. Smith, and L. Kong, “Random feature attention,” in *Presented at the International Conference on Learning Representations*, 2020.
- [10] R. Sato, M. Yamada, and H. Kashima, “Random features strengthen graph neural networks,” in *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*, in *Proceedings*, Society for Industrial and Applied Mathematics, 2021, pp. 333–341.
- [11] A. Saade, *et al.*, “Random projections through multiple optical scattering: approximating kernels at the speed of light,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 6215–6219.
- [12] U. Teğin, M. Yildirim, İ. Oğuz, C. Moser, and D. Psaltis, “Scalable optical learning operator,” *Nat. Comput. Sci.*, vol. 1, no. 8, pp. 542–549, 2021.
- [13] L. G. Wright, *et al.*, “Deep physical neural networks trained with backpropagation,” *Nature*, vol. 601, no. 7894, 2022, Art. no. 7894. <https://doi.org/10.1038/s41586-021-04223-6>.
- [14] M. Yildirim, *et al.*, “Nonlinear optical feature generator for machine learning,” *APL Photonics*, vol. 8, no. 10, 2023, Art. no. 106104. <https://doi.org/10.1063/5.0158611>.
- [15] L. G. Wright, D. N. Christodoulides, and F. W. Wise, “Controllable spatiotemporal nonlinear effects in multimode fibres,” *Nat. Photonics*, vol. 9, no. 5, pp. 306–310, 2015.
- [16] B. U. Kesgin and U. Teğin, “Photonic neural networks with spatiotemporal chaos in multimode fibers,” *arXiv*: arXiv:2411.00189, 2024. <https://doi.org/10.48550/arXiv.2411.00189>.
- [17] R. Shao, G. Zhang, G. Zhang, X. Gong, and X. Gong, “Generalized robust training scheme using genetic algorithm for optical neural networks with imprecise components,” *Photonics Res.*, vol. 10, no. 8, pp. 1868–1876, 2022.
- [18] I. Oguz, *et al.*, “Programming nonlinear propagation for efficient optical learning machines,” *Adv. Photonics*, vol. 6, no. 1, 2024, Art. no. 016002. <https://doi.org/10.1117/1.AP.6.1.016002>.
- [19] I. Oguz, *et al.*, “Forward—forward training of an optical neural network,” *Opt. Lett.*, vol. 48, no. 20, pp. 5249–5252, 2023.
- [20] D. Wang, *et al.*, “Digital twin of optical networks: a review of recent advances and future trends,” *J. Lightwave Technol.*, vol. 42, no. 12, pp. 4233–4259, 2024.
- [21] O. Bernard and Y. Bellouard, “On the use of a digital twin to enhance femtosecond laser inscription of arbitrary phase patterns,” *J. Phys.: Photonics*, vol. 3, no. 3, 2021, Art. no. 035003. <https://doi.org/10.1088/2515-7647/abf743>.
- [22] N. Borhani, E. Kakkava, C. Moser, and D. Psaltis, “Learning to see through multimode fibers,” *Optica*, vol. 5, no. 8, p. 960, 2018.
- [23] B. Rahmani, *et al.*, “Actor neural networks for the robust control of partially measured nonlinear systems showcased for image propagation through diffuse media,” *Nat. Mach. Intell.*, vol. 2, no. 7, pp. 403–410, 2020.
- [24] A. Antolini, *et al.*, “Combined HW/SW drift and variability mitigation for PCM-based analog in-memory computing for neural network applications,” *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 13, no. 1, pp. 395–407, 2023.

- [25] M. Hasan, C. Nicholls, K. Pitre, B. Spokoinyi, and T. Hall, “Delay drift compensation of an optoelectronic oscillator over a large temperature range through continuous tuning,” *Commun. Eng.*, vol. 3, no. 1, pp. 1–8, 2024.
- [26] H. Cao, T. Čížmár, S. Turtaev, T. Tyc, and S. Rotter, “Controlling light propagation in multimode fibers for imaging, spectroscopy, and beyond,” *Adv. Opt. Photonics*, vol. 15, no. 2, pp. 524–612, 2023.
- [27] K. Wisal, C.-W. Chen, Z. Kuang, O. D. Miller, H. Cao, and A. D. Stone, “Optimal input excitations for suppressing nonlinear instabilities in multimode fibers,” *Optica*, vol. 11, no. 12, pp. 1663–1672, 2024.
- [28] H. Xiao, K. Rasul, and R. Vollgraf, “Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms,” *ArXiv170807747 Cs Stat*, 2017. <http://arxiv.org/abs/1708.07747>.
- [29] S. Biasi, R. Franchi, L. Cerini, and L. Pavesi, “An array of microresonators as a photonic extreme learning machine,” *APL Photonics*, vol. 8, no. 9, 2023, Art. no. 096105. <https://doi.org/10.1063/5.0156189>.
- [30] X. Lin, *et al.*, “All-optical machine learning using diffractive deep neural networks,” *Science*, vol. 361, no. 6406, pp. 1004–1008, 2018.
- [31] E. Kazakov, J. Gao, P. Anisimov, and V. Zemlyakov, “Parallelization of the generalized multimode nonlinear schrödinger equation solver: a performance analysis,” in *Parallel Computational Technologies*, L. Sokolinsky, and M. Zymbler, Eds., Cham, Springer Nature Switzerland, 2023, pp. 137–151.
- [32] O. Ronneberger, P. Fischer, and T. Brox, “U-net: convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, Springer, 2015, pp. 234–241.
- [33] J. P. Kopolow, D. A. V. Kliner, and L. Goldberg, “Single-mode operation of a coiled multimode fiber amplifier,” *Opt. Lett.*, vol. 25, no. 7, pp. 442–444, 2000.
- [34] R. W. Boyd, *Nonlinear Optics*, 3rd ed. USA, Academic Press, Inc., 2008.
- [35] G. Agrawal, “Nonlinear science at the dawn of the 21st century,” in *Nonlinear Fiber Optics*, Berlin, Heidelberg, Springer Berlin Heidelberg, 2000, pp. 195–211.
- [36] A. Mafi, “Pulse propagation in a short nonlinear graded-index multimode optical fiber,” *J. Lightwave Technol.*, vol. 30, no. 17, pp. 2803–2811, 2012.