

Quantum optical neural networks using atom-cavity interactions to provide all-optical nonlinearity

Chuanzhou Zhu,^{1,*} Tianyu Wang,² Peter L. McMahon,³ and Daniel Soh^{1,†}

¹*Wyant College of Optical Sciences, University of Arizona, Tucson, Arizona, USA*

²*Department of Electrical & Computer Engineering, Boston University, Boston, Massachusetts, USA*

³*School of Applied and Engineering Physics, Cornell University, Ithaca, New York, USA*

(Dated: November 11, 2025)

Optical neural networks (ONNs) have been developed to enhance processing speed and energy efficiency in machine learning by leveraging optical devices for nonlinear activation and establishing connections among neurons. In this work, we propose a quantum optical neural network (QONN) that utilizes atom-cavity neurons with controllable photon absorption and emission. These quantum neurons are designed to replace the electronic components in ONNs, which typically introduce delays and substantial energy consumption during nonlinear activation. To evaluate the performance of the QONN, we apply it to the MNIST digit classification task, considering the effects of photon absorption duration, random atom-cavity detuning, and stochastic photon loss. Additionally, we introduce a convolutional QONN to facilitate a real-world satellite image classification (SAT-6) task. Due to its compact hardware and low power consumption, the QONN offers a promising solution for real-time satellite sensing, reducing communication bandwidth with ground stations and thereby enhancing data security.

I. INTRODUCTION

Recent advancements in artificial intelligence have been driven by the rapid growth in the number of internal parameters, such as the trillions of parameters in models like GPT-5. This expansion has led to a substantial increase in both energy consumption and the demand for faster processing speeds [1, 2].

Optical neural networks (ONNs) have been proposed as a potential solution to enhance processing speed and carrier frequency, benefiting from the broader optical bandwidth (typically in the terahertz range) compared to the electronic bandwidth (typically in the gigahertz range) [3–20]. ONNs have also achieved better energy efficiency than traditional electronic processors. In typical low-energy ONNs, millions of photons are necessary for the neuron pre-activation signal to ensure reliable performance [9, 10, 13, 17]. Recent advancements in ONN energy efficiency have reduced photon usage to hundreds of photons per activation [11]. In the regime of ultra-low photon numbers, stochastic optical neural networks have been implemented to perform accurate machine learning inference, despite the challenges posed by low signal-to-noise ratios [21]. The energy efficiency in these implementations is often measured by the number of photons involved in optical matrix-vector multiplications, which are performed through parallel vector-vector dot products enabled by techniques such as wavelength multiplexing [8, 9, 22], spatial multiplexing in photonic integrated circuits [4, 23–25], and spatial multiplexing in 3D free-space optical processors [26–33].

However, two key sources of speed delay and energy consumption remain largely unaddressed in ONNs: the

nonlinear activation following a multiplication operation and the photon emission required before the next multiplication stage. Both processes rely on macroscopic electronic devices, such as single-photon detectors and photon emitters [11, 21], which lead to the electronic delay in processing speed and contribute significantly to overall energy consumption at the macroscopic level.

Quantum neural networks (QNNs) [34–47] have been proposed as quantum generalizations of classical neural networks, implemented on small-scale quantum platforms such as superconducting quantum processors [42, 44]. These architectures might provide benefits due to uniquely quantum phenomena such as entanglement. A QNN comprises a set of parameterized quantum logical gates that perform nonlinear activation in each neuron and produce quantum entanglement across multiple neurons. To achieve the desired outputs from QNNs, the gate parameters are typically trained using classical [48] and quantum [34] optimizers. However, quantum decoherence increases the complexity of designing reliable quantum gates [49, 50], particularly for two-qubit gates [51, 52], which has spurred efforts to develop fault-tolerant quantum gates through topological protection [53, 54] and error correction [55]. Additionally, the need for fine-tuning internal gate parameters further complicates the fabrication of robust QNNs.

Quantum optical neural networks (QONNs) [56–59] have emerged as novel platforms for implementing linear transformations and nonlinear activation functions in neural networks through quantum optical systems. Although these QONNs do not rely directly on qubit gates as in QNNs, they still employ gate components such as continuous-variable gates [59] and optical unitaries [56, 57], along with internal gate parameters that must be precisely adjusted. Actually, even before machine learning and artificial intelligence entered the public consciousness, quantum-optics-based neural networks

* chuanzhouzhu@arizona.edu

† danielsoh@arizona.edu

had already been proposed to simulate information processing in neurobiological experiments. Early examples include a network composed of two neural states exhibiting transitions analogous to photon-emission models in quantum optics [60, 61], a quantum network hardware architecture constructed from a spatial array of quantum dot molecules [62], and an optical neural network built with quantum well devices employing two arrays of integrated asymmetric Fabry-Perot modulators [63]. More recently, a QONN has been proposed by mapping the features of modern machine-learning neural networks onto the few-photon quantum optical domain [56]. In this framework, the nonlinear activation is realized through a phase quadratic generated via single-mode Kerr interactions, while the matrix multiplication is implemented using a linear optical unitary achieved through arrays of beam splitters and programmable phase shifts [64]. Subsequent studies have investigated the effects of photon propagation loss, weak nonlinearities, and network size on QONN performance [57]. Furthermore, it has been proposed that nonlinear activation in QONNs can also be produced using the Hong-Ou-Mandel effect [58]. However, in QONNs that rely on the Kerr effect and optical unitaries, the performance is constrained by the weak Kerr nonlinearity in the few-photon regime [65]. Moreover, tuning the trainable parameters in the optical unitary does not enable modulation of the average photon amplitude.

We present a QONN that employs atom-cavity neurons to replace the electronic photon detectors and emitters in ONNs, while avoiding the use of quantum gates or the adjustment of internal gate parameters as seen in QNNs [34–47] and previous proposals for QONNs [56–59]. Our QONN harnesses cavity arrays and optical matrix-vector multipliers (MVMs) to transmit single-photon amplitudes, which serve as information carriers across multiple layers of the network. Each neuron in the cavity array performs nonlinear activation on the optical signal and is constructed with an atom placed inside two cavities, which can be switched on and off to absorb and emit photons. Our results show that the desired nonlinearity in the activation function can be achieved by tuning the photon absorption duration. The optical MVM is implemented via a controllable spatial light modulator (SLM) that performs fully connected transformations or convolution operations. Experimental results have demonstrated that flexible parameter adjustment in optical MVMs enables high inference accuracy in ONNs [11, 21].

The processing speed in our setup is primarily determined by the Rabi frequency of the cavity neuron, which typically falls within the megahertz range in conventional atom-photon coupling experiments [66, 67]. In contrast, all-optical setups involving photonic qubits can achieve Rabi frequencies in the terahertz range [68–70], offering the potential to substantially accelerate processing speeds. The all-optical implementation of our cavity quantum electrodynamics neurons can be realized using

quantum dots embedded in a photonic crystal [71–73]. While several photon detectors remain necessary at the output layers, our setup eliminates the need for thousands of single-photon detectors and emitters typically required in the hidden layers in previous ONN experiments [11, 21], thereby significantly reducing the overall energy consumption. Due to its compact hardware and low energy consumption, the QONN is suitable as an onboard learning system for real-time remote sensing on satellites, reducing the need for high-bandwidth communication with ground stations and thereby improving data security.

Photon entanglement is not included in the present model, as we use a mean-field treatment for the photon field. Under realistic conditions, photon entanglement naturally arises in optical MVMs, potentially enhancing the capabilities of the QONN by exploiting a larger Hilbert space. The performance enhancements related to quantum entanglement have been studied in other quantum artificial intelligence frameworks, such as quantum reservoir computing [74–76]. QONN models based on continuous-variable gates and optical unitaries [56–59] are well suited to simulate quantum features associated with entanglement.

To demonstrate the capability of QONN in image recognition, we adopt MNIST handwritten digit classification [77] as a benchmark task to identify the optimal photon absorption duration for each cavity array that performs as a hidden layer. This is followed by simulations under realistic conditions, incorporating random atom-cavity detuning in each cavity neuron and stochastic single-photon loss during photon transmission. To highlight the real-world applicability of QONN in remote sensing, we test the benchmarked system on a satellite image classification task, namely the DeepSat (SAT-6) airborne image classification [78]. To further optimize the system, we propose a convolutional QONN to reduce the number of controllable SLM pixels without compromising accuracy or requiring additional cavity neurons.

II. SETUP

The schematic of the QONN for image recognition is illustrated in Fig. 1a. It consists of an input layer, two fully connected hidden layers, and an output layer, where the layers are indexed by l and the neurons in each layer by i . The framework can be extended to include additional hidden layers to further enhance performance. Each pixel in an input image is encoded as a photon amplitude, denoted by a_i^{in} . The optical matrix-vector multiplier (MVM), implemented via a spatial light modulator (SLM), linearly connects the activation function a_i^{l-1} of the $(l-1)$ -th layer to the incident photon amplitude z_i^l of the l -th layer. The controllable transmission of each SLM pixel realizes the trainable parameter W_{ij}^l in the network. The cavity array, composed of decoupled cavity neurons, performs quantum optical activation by




FIG. 1. Quantum optical neural network for MNIST digit classification and DeepSat (SAT-6) airborne image classification tasks. **a** Schematic of the quantum optical neural network, consisting of optical matrix-vector multipliers (MVMs) and cavity arrays, where z_i^l represents the incident photon amplitude entering each cavity neuron, and a_i^l is the photon amplitude emitted by the atom inside the cavity neuron. The activation values a_i^l from the cavity array of each layer are passed to the next layer through the optical MVM, with z_i^{l+1} being the weighted sum of a_i^l . The optical MVM is implemented using a spatial light modulator (SLM). The weights W_{ij}^l , which are parameters optimized via backpropagation, are adjusted by controlling the transmission of each SLM pixel. The input images are encoded in a_i^{in} . The final layer uses a photodetector (PD) array to obtain the output a_i^{out} , which is used to compute the cost function. **b** Schematic of a cavity neuron that performs the quantum optical nonlinear activation as described by Eq. (7). The neuron consists of a low-Q cavity with weaker atom-cavity coupling (in blue) and a high-Q cavity with stronger atom-cavity coupling (in red). In step 1, the incident photon is absorbed by the two-level atom, and the excitation is stored in the low-Q cavity while the high-Q cavity is turned off. In step 2, the high-Q cavity is turned on, and the excitation is transferred from the low-Q cavity to the high-Q cavity. In step 3, the high-Q cavity is re-opened, and the photon is emitted by the atom through a complete energy conversion from atomic excitation to photon emission.

establishing a nonlinear relation between each incident photon amplitude z_i^l and the emitted photon amplitude a_i^l . The photodetector (PD) array in the final layer retrieves the output photon amplitude a_i^{out} , which corresponds to the inference result and is utilized to calculate the cost function.

Quantum nonlinear activation is realized through the cavity neuron illustrated in Fig. 1b. Each neuron comprises a two-level atom surrounded by two optical cavities: one with a lower Q factor and weaker atom-cavity coupling strength (blue), and another with a higher Q factor and stronger atom-cavity coupling (red). The high-Q cavity can be switched on or off by suppressing or enhancing photon loss using a phase shifter, which controls the destructive interference between the leftward and rightward traveling waves of the cavity mirror system [79, 80]. The nonlinear activation process involves three steps: photon absorption, excitation transfer, and spontaneous emission.

Step 1: The high-Q cavity is turned off, and the low-Q cavity facilitates photon absorption by the two-level atom. The incident photon acts as an external drive for the atom. For the i -th cavity neuron in the l -th layer, the absorption process is governed by the Hamiltonian

$$H_i^l = \delta_i^l \sigma_z + g (\sigma_+ z_i^l + \sigma_- z_i^{l*}), \quad (1)$$

where z_i^l is the amplitude of the incident photon field under the mean-field approximation, δ_i^l denotes the atom-cavity detuning, g represents the atom-photon electric

dipole coupling strength, and σ_z and $\sigma_{\pm}$ are Pauli matrices. Each atom is initialized in the ground state, and the atomic excitation follows a Rabi oscillation described by

$$\langle \sigma_z \rangle_i^l = - \left(\frac{\delta_i^l}{\Omega_i^l} \right)^2 - \left(\frac{gz_i^l}{\Omega_i^l} \right)^2 \cos(2\pi t_l \Omega_i^l), \quad (2)$$

where the oscillation frequency depends on z_i^l as

$$\Omega_i^l(z_i^l) \equiv \sqrt{(gz_i^l)^2 + (\delta_i^l)^2}, \quad (3)$$

and t_l is the absorption time, which is assumed identical for all neurons in the same layer.

Step 2: The high-Q cavity is turned on. Since the atom tends to more rapidly interact with the high-Q cavity through its photon absorption and emission within the two-cavity neuron, the excitation is transferred from the low-Q cavity to the high-Q one.

Step 3: Once all cavity neurons in a layer are excited, all high-Q cavities in that layer are simultaneously opened. Through complete energy conversion from atomic excitation to photons, the intensity of the spontaneously emitted photon field is given by

$$|a_i^l|^2 = \frac{1}{2} (\langle \sigma_z \rangle_i^l + 1), \quad (4)$$

where a_i^l represents the emitted photon amplitude. Accounting for a random phase ϕ_i^l induced by spontaneous emission, the emitted photon amplitude is expressed as

$$a_i^l = |a_i^l| e^{i\phi_i^l}, \quad (5)$$

where $|a_i^l|$ follows Eq. (4) and $\langle \sigma_z \rangle_i^l$ is defined in Eq. (2).

The activation values a_i^l are linearly propagated to the next layer through optical MVM, implemented by displaying these values on the SLM, as shown in Fig. 1a. The SLM consists of a pixel array with individually controllable transmission for each pixel, as demonstrated in Refs. [11, 21]. Element-wise multiplication is achieved through amplitude modulation, resulting in the linear combination

$$z_i^l = \sum_j W_{ij}^l a_j^{l-1}, \quad (6)$$

where the weight W_{ij}^l is tuned by adjusting the transmission of a single SLM pixel, and hence constrained within $W_{ij}^l \in [-1, 1]$. These weights serve as trainable parameters of the neural network, which are optimized during the training process described in Appendix B.

To eliminate the randomness of the phase ϕ_i^l in the activation function (5), a weak auxiliary external laser need to be applied to induce stimulated emission for phase locking [81, 82]. The locked phase ϕ_i^l can then be compensated by the complex phase of the transmission rate W_{ij}^l , allowing a_i^l , z_i^l , and W_{ij}^l to be treated as real-valued. Consequently, the auxiliary phase-locking process yields the nonlinear quantum activation function

$$a_i^l(z_i^l) = \frac{g|z_i^l|}{\Omega_i^l(z_i^l)} |\sin[\pi t_l \Omega_i^l(z_i^l)]|, \quad (7)$$

where $\Omega_i^l(z_i^l)$ is a nonlinear function of z_i^l defined by Eq. (3). The pre-activation and activation values span $z_i^l \in (-\infty, +\infty)$ and $a_i^l \in [0, 1]$, respectively.

The design of our two-cavity neuron supports temporal multiplexing for optically decomposing the matrix-vector multiplication in Eq. (6) into a batch of vector-vector dot products, which are executed in time sequence. For instance, the input photon amplitudes a_i^{in} are transmitted through the SLM pixels and focused onto the first neuron in the first hidden layer, leading to the vector-vector dot product $z_1^1 = \sum_j W_{1j}^1 a_j^{\text{in}}$. The excitation is then held during *Step 2* by maintaining the high-Q cavity (red) in the “on” state, preserving the excitation inside the high-Q cavity without photon emission. This procedure sequentially excites all neurons in the first hidden layer, corresponding to $z_i^1 = \sum_j W_{ij}^1 a_j^{\text{in}}$ for each of i -th neurons in the layer 1. After the entire first hidden layer is excited, all neurons in this layer simultaneously proceed to *Step 3*, during which the high-Q cavities are opened and all emitted photons are transmitted to the first neuron in the second hidden layer, producing $z_1^2 = \sum_j W_{1j}^2 a_j^1$. This temporal multiplexing continues throughout the network.

The output layer employs a PD array to measure the amplitudes z_i^{out} , with normalized measurements represented as a_i^{out} . The number of photodetectors corresponds to the number of categories in the specific classification task. In the MNIST digit recognition task, the target output is defined as $a_{i=\text{digit}}^{\text{out}} = 1$ and $a_{i \neq \text{digit}}^{\text{out}} = 0$,

where *digit* denotes the input handwritten numeral. The deviation from the ground truth is quantified by a cost function, where the LogSoftmax function is applied to z_i^{out} and the cross-entropy loss is used for optimization. The neural network is trained using backpropagation with the stochastic gradient descent (SGD) optimizer. Further details regarding the training and testing procedures are provided in Appendix B.

III. PERFORMANCE EVALUATION UNDER REALISTIC CONDITIONS

A. Photon Absorption Durations

According to the formalism of the activation function in Eq. (7), the photon absorption time t_l in *Step 1* of the quantum activation determines the degree of nonlinearity and monotonicity between the pre-activation z_i^l and the activation value a_i^l . When the detuning $\delta_i^l = 0$, the activation function simplifies to

$$a_i^l(z_i^l) = |\sin(\pi t_l g |z_i^l|)|. \quad (8)$$

Based on the setup with two hidden layers, Fig. 2a illustrates the test accuracy for MNIST digit classification as a function of the photon absorption times in the first and second hidden layers, denoted by t_1 and t_2 , respectively. When $t_1 = t_2 = t$, Fig. 2b presents the improvement of test accuracy as the number of training epochs increases (i.e., repeated iterations over the MNIST dataset). The performance tends to saturate after 10 epochs. The dependence of accuracy on t arises from the trade-off between the nonlinearity and monotonicity of the activation function. As shown in Fig. 2c, for a short absorption duration ($t = 0.2$), most pre-activation values z_i^l fall within the nearly linear regime of the activation function, rendering the activation insufficiently nonlinear to achieve high accuracy. Conversely, for a long absorption duration ($t = 4$), the rapid oscillations of the activation function disrupt the monotonic correspondence between z_i^l and a_i^l , such that different incident amplitudes produce identical photon emissions. This loss of one-to-one mapping leads to information degradation and reduced accuracy. The optimal performance occurs around an intermediate absorption time ($t = 1$), where the nonlinearity and monotonicity are effectively balanced.

The photon absorption durations in the first and second hidden layers, t_1 and t_2 , play distinct roles in determining accuracy. The asymmetric colormap in Fig. 1a indicates that optimal performance requires tighter control over t_1 , whereas a broader range of t_2 values yields acceptable accuracies. This asymmetry can be understood from the parameter training process: once the activation functions in both layers are fixed by a given combination of t_1 and t_2 , the training process can only adjust the weights W_{ij}^l to optimize the distribution of neuron populations in the z_i^l domain, thereby balancing the benefits




FIG. 2. Performance of the quantum optical neural network with various photon absorption times for the MNIST task. **a** Test accuracy plotted against the photon absorption times in the first and second hidden layers, t_1 and t_2 . **b** Test accuracy as a function of the absorption time t , with $t = t_1 = t_2$, after 1 (green), 10 (red), and 20 (orange) training epochs. A training epoch is defined as a complete pass of the entire MNIST training dataset through the learning process. **c** Nonlinear activation function a_i^l (identical for all neurons in the two layers since $t_1 = t_2$), and the occurrence population distributions of the incident photon amplitudes $|z_i^1|$ and $|z_i^2|$, which serve as the arguments for the activation functions. The results shown in panel **c** correspond to three absorption times, $t = 0.2$, $t = 1$, and $t = 4$, based on 10 training epochs and the input image shown in the upper-left corner. The detuning is $\delta_i^l = 0$, and the coupling strength is $g = 1$ for every cavity neuron. The number of neurons is set to $N_1 = N_2 = 512$ in the two hidden layers.

of nonlinearity and monotonicity in the first layer. In contrast, for the second layer, the training process can further tune both W_{ij}^1 and W_{ij}^2 to reach the optimal population distribution of z_i^2 values. This interpretation is consistent with the distinct distributions of z_i^1 and z_i^2 illustrated in Fig. 2c.

B. Random Detuning

In all-optical implementations using quantum dots embedded in photonic crystals to realize artificial two-level atoms inside optical cavities, the atom-photon detuning in each cavity is typically difficult to control and is often modeled as a random variable [71–73, 83]. Figure 3a presents the influence of such random detuning on the MNIST test accuracy, where the activation function of each neuron is simulated using Eq. (7) with a finite detuning δ_i^l . The detuning δ_i^l for each neuron is independently drawn from a uniform probability distribution over the range between $-2\delta_0$ and $2\delta_0$. The QONN exhibits robust

performance against a wide range of random detunings. It is worth noting that the values of random detuning shown in Fig. 3a (up to $\delta_0 = 2.5$) represent substantial magnitudes, given that the coupling strength is set to $g = 1$ in the simulation and the typical magnitude of $|z_i^l|$ is on the order of 1 as evidenced in Fig. 2c.

For comparison, Fig. 3b presents the case where all cavity neurons are identical and share the same detuning $\delta_i^l = \delta_0$. In this scenario, the network achieves better accuracy only for a few specific combinations of δ_0 and t , corresponding to occasional optimal balances between nonlinearity and monotonicity in the activation function. In contrast, random detuning relaxes the requirements for precise control over both absorption time and detuning. The generally higher accuracy achieved under random detuning highlights the advantage of introducing variability among cavity neurons, enabling them to develop non-identical response characteristics and thereby enhancing the collective computational capability of the network.




FIG. 3. Performance of the quantum optical neural network with finite detuning for each cavity neuron on the MNIST task. The photon absorption times t are identical in the first and second hidden layers. **a** Each cavity neuron is simulated with a random detuning δ_i^l , where the probability is uniformly distributed within the range $[-2\delta_0, 2\delta_0]$, resulting in an average detuning magnitude of δ_0 . **b** All cavity neurons are identical and simulated with a fixed detuning $\delta_i^l = \delta_0$. The coupling strength is $g = 1$, and the number of neurons is $N_1 = N_2 = 512$.

C. Stochastic Photon Loss

Stochastic photon loss is an inherent and practically unavoidable phenomenon in optical transmission systems [84, 85]. Figure 4a illustrates a stochastic QONN model to simulate the photon loss, where a stochastic layer is inserted before the emitted photons enter the optical MVM stage of the next layer. To simulate the single-photon loss semi-classically, we assume that each photon has a probability P of passing through the stochastic layer unperturbed and a probability $1 - P$ of being completely lost to the environment. This stochastic loss occurs independently for the photons emitted from different cavity neurons, with an identical transmission probability P for all photons. During both the training and testing stages, the stochastic layer is included in the forward propagation. In the backward propagation during training, however, a mean-field approximation is employed, treating photon loss deterministically, as detailed in Appendix C. The impact of stochastic photon loss on the MNIST classification task is presented in Fig. 4b, where a single hidden layer is considered for simplicity. Remarkably, the test accuracy remains around 80% even when the photon pass rate decreases to 20%, demonstrating the strong robustness of the QONN against photon loss when performing classification tasks.

D. SAT-6 Task and Convolutional Neural Network

The DeepSat (SAT-6) airborne image classification task [78] is employed to demonstrate a real-world application of the QONN. The task is introduced in Appendix




FIG. 4. Quantum optical neural network with stochastic photon loss. **a** Based on the schematic shown in Fig. 1a, a stochastic layer (in magenta) is added before the activation a_i^1 entering the optical MVM to model photon loss during transmission. During the forward pass, each a_i^1 has a probability P of remaining unchanged and a probability $1 - P$ of being set to 0. The photon passing rate P is identical for all cavity neurons. **b** MNIST test accuracy plotted against the photon passing rate. The stochastic neural network consists of a single hidden layer with 512 cavity neurons, each having a photon absorption time $t_1 = 1$ and a detuning $\delta_i^1 = 0$, and is trained for 10 epochs.

A. Owing to its small hardware size and low-energy consumption, the QONN can serve as an on-board learning machine for real-time remote sensing on satellites. It reduces the communication bandwidth required between the satellite and ground stations, thereby enhancing both operational efficiency and data security. The task involves classifying satellite images into six land-cover categories, as illustrated in Fig. 5a.

To efficiently process images with RGB color channels, an optical convolutional layer is introduced, as shown in Fig. 5b, to reduce the number of trainable parameters. Unlike a fully connected (FC) layer, where all emitted




FIG. 5. Quantum optical neural network performing DeepSat (SAT-6) airborne image classification task. **a** SAT-6 images with RGB channels representing six land cover classes. **b** Schematic of a convolutional layer embedded in the quantum optical neural network, where the convolution kernel is implemented using a set of SLMs, followed by average pooling realized through programmable photonic circuits. The photons are incident into the cavity neurons after the convolution and pooling operations. **c** Test accuracy of the quantum optical neural network on the SAT-6 task with varying numbers of cavity neurons in the first hidden layer. Two network structures are applied: one with a convolutional layer followed by a fully-connected layer (orange), and another with two fully-connected layers as shown in Fig. 1a (blue). The number of parameters in the convolutional layer depends on the kernel size and the number of output channels. A 5×5 kernel with a stride of 1 and 2×2 average pooling with a stride of 2 are used. The number of output channels in the convolutional layer is set to 2, 4, 8, and 16, corresponding to the four orange dots in panel c, respectively. The second layer contains 512 neurons in both structures. The network is trained for 10 epochs.

photons are focused onto a single cavity, in the convolutional layer each cavity processes photons originating from a small spatial region (e.g., a 5×5 kernel in Fig. 5b). The same sets of convolutional filters are applied across multiple regions to extract shared features. The convolutional filter is performed by optical MVM with multiple SLMs, with the adjustable transmission rate of each pixel on SLMs treated as a parameter in the convolutional QONN. Three SLMs are encompassed in each set of filters to address the three color channels, and the number of the filter sets corresponds to the number of output channels. The average pooling, implemented by programmable photonic circuits [24], combines the photon amplitudes after the convolution operation. After the convolution and pooling, photons are passed to the

cavity array for nonlinear activations.

Figure 5c plots the SAT-6 test result obtained using a convolutional QONN (in orange), which adopts the two-layer structure shown in Fig. 1a but replaces the first hidden layer with the convolutional layer depicted in Fig. 5b. For comparison, the original structure with two FC layers is also used to perform SAT-6 task using the same number of neurons (in blue). In the convolutional network, the number of parameters is determined by the kernel size and the number of output channels, with the details discussed in Appendix D. The FC network, however, requires a significantly larger number of parameters, which is decided by the product of the neuron counts in the two connected layers. The results in Fig. 5c highlight the advantage of the convolutional QONN, which substantially

reduces the control complexity of SLM pixel transmission rates while maintaining comparable test accuracy without the need for additional cavity neurons.

IV. DISCUSSION

We have proposed a QONN to address the electronic delay and energy consumption issues in ONNs without the need for complex qubit gate design as seen in QNNs. The QONN harnesses quantum optical neurons to produce nonlinear activations through absorbing and emitting single photons, supporting both fully-connected and convolutional layers. The effects of photon absorption duration, random atom-photon detuning, and stochastic photon loss have been investigated. The QONN has achieved over 95% accuracy on both the benchmark MNIST digit classification and the real-world SAT-6 task.

Quantum entanglement can expand the dimensionality of the computational Hilbert space spanned by a large number of basis states, which has been shown to enhance the capabilities of some other quantum artificial intelligence systems, such as quantum reservoir computing [75, 76]. In the QONN, photon-photon entanglement occurs in the optical MVM when multiple photon amplitudes are combined as a pre-activation. This entanglement is not accounted for in our current mean-field treatment of the incident and activation photon fields, as an analytical formalism for the activation function, given by Eq. 7, is required for gradient-based training. A fully quantum treatment beyond the mean-field approximation represents an important future research direction, which could demonstrate the quantum advantage of QONNs and enable comparisons with classical counterparts. Alternatively, within the current framework, a larger computational Hilbert space could be constructed by connecting the cavity neurons using optical waveguides or by adding more atoms to each cavity. These investigations could also lead to the discovery of more efficient quantum activation functions through balancing nonlinearity and monotonicity.

ACKNOWLEDGEMENTS

This work is supported by the U.S. Department of Energy, Office of Science, Office of Biological and Environmental Research under Award Number DE-SCO025910.

Appendix A: MNIST and SAT-6 Tasks

The benchmark MNIST (Modified National Institute of Standards and Technology) handwritten digit dataset comprises 60000 training images and 10000 test images [77]. Each image is a 28×28 grayscale matrix, resulting in 784 pixels per image. To encode the image onto the

input photon amplitude, all pixel values are normalized to fall within the range $[0, 1]$.

The real-world DeepSat (SAT-6) airborne image classification task consists of 324000 training images and 81000 test images, which fall into six land cover classes including buildings, barren land, trees, grassland, roads, and water bodies [78]. Each image consists of 28×28 pixels and 4 color channels - red, green, blue (RGB) and near infrared. We only use the RGB channels, leading to 2352 data points per image.

Appendix B: Backward Propagation for Training

The cost function $C(a_i^{\text{out}})$ characterizes the error of the inference by comparing the output a_i^{out} with the ground truth. We employ the cross-entropy loss method [86], which applies LogSoftmax function on z_i^{out} to simulate the nonlinearity induced by the photodetector array in the output layer, followed by the negative log likelihood loss (NLLLoss) [87] to evaluate the cost function. The backward propagation is processed with chain rule as

$$\frac{\partial C}{\partial z_i^l} = \frac{\partial a_i^l}{\partial z_i^l} \frac{\partial C}{\partial a_i^l}, \quad (\text{B1})$$

$$\frac{\partial C}{\partial a_i^{l-1}} = \sum_j \frac{\partial C}{\partial z_j^l} W_{ji}^l, \quad (\text{B2})$$

where $\partial a_i^l / \partial z_i^l$ is analytically expressed as the derivative of the activation function $a_i^l(z_i^l)$ given by Eq. (7). For example, the layers in Fig. 1a are processed backward in the order of $l = \text{out}, 2, 1, \text{in}$. The gradient of the cost function with respect to the weight W_{ij}^l is computed as

$$g_{W_{ij}^l} \equiv \frac{\partial C}{\partial W_{ij}^l} = \frac{\partial C}{\partial z_i^l} a_j^{l-1}, \quad (\text{B3})$$

where $\partial C / \partial z_i^l$ is resulted from Eq. (B1). After running through a batch of training images, we use stochastic gradient descent (SGD) optimizer to update the weights based on the gradient values obtained from Eq. (B3). During the training, we restrict the value of W_{ij}^l in the range $[-1, 1]$ in count of the actual photon transmission rate of a SLM pixel. We set the batch size to 64 and the learning rate to 10^{-3} , and use PyTorch in the simulation.

Appendix C: Stochastic Layer

The photon pass rate P is introduced to characterize the single-photon loss in the stochastic layer. As shown in Fig. 4a, in the forward propagation for training and testing, the stochastic photon loss is simulated as

$$a_i^{\text{ls}} = f(P) a_i^1, \quad (\text{C1})$$

where $f(P)$ is a random variable with a Bernoulli distribution: $f(P) = 1$ with probability P , and $f(P) = 0$

with probability $1 - P$. In the backward propagation for training, the deterministic mean-field approximation is implemented by assuming

$$a_i^{1s} = P a_i^1. \quad (\text{C2})$$

This correspondingly modifies the gradient in the stochastic QONN as

$$g_{W_{ij}^l}^{\text{stochastic}} = P g_{W_{ij}^l}, \quad (\text{C3})$$

where $g_{W_{ij}^l}$ is defined in Eq. (B3). The modified gradient is involved in the weights update through the SGD optimizer for both layers indexed by $l = \text{out}$ and $l = 1$ in Fig. 4a.

Appendix D: Convolutional Layer

In the convolutional layer illustrated in Fig. 5b, the input SAT-6 image offers a feature map of $28 \times 28 \times 3$.

It is followed by a convolution operation with a kernel size of 5×5 , a stride size of 1, a padding of 0, and N_{channels} output channels, which results in a data dimension of $24 \times 24 \times N_{\text{channels}}$. After that, an average pooling of 2×2 is applied to reduce the data dimension to $12 \times 12 \times N_{\text{channels}}$, which is then flattened into a vector and applied to the neurons in the first layer of the network. In Fig. 5c, we consider $N_{\text{channels}} = 2, 4, 8, 16$, which necessitates the number of neurons from 288 to 2304. The number of parameters in the convolutional layer is given by $5 \times 5 \times 3 \times N_{\text{channels}}$ with the 5×5 kernel and 3 input color channels. This number varies from 150 to 1200 in Fig. 5c.

-
- [1] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature* **521**, 436 (2015).
 - [2] N. C. Thompson, K. Greenewald, K. Lee, and G. F. Manso, *The computational limits of deep learning* (2022), [arXiv:2007.05558 \[cs.LG\]](#).
 - [3] N. L. Kazanskiy, M. A. Butt, and S. N. Khonina, Optical computing: Status and perspectives, *Nanomaterials* **12**, 10.3390/nano12132171 (2022).
 - [4] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and M. Soljačić, Deep learning with coherent nanophotonic circuits, *Nature Photonics* **11**, 441 (2017).
 - [5] X. Lin, Y. Rivenson, N. T. Yardimci, M. Veli, Y. Luo, M. Jarrahi, and A. Ozcan, All-optical machine learning using diffractive deep neural networks, *Science* **361**, 1004 (2018).
 - [6] C. Rios, N. Youngblood, Z. Cheng, M. L. Gallo, W. H. P. Pernice, C. D. Wright, A. Sebastian, and H. Bhaskaran, In-memory computing on a photonic platform, *Science Advances* **5**, eaau5759 (2019).
 - [7] G. Wetzstein, A. Ozcan, S. Gigan, S. Fan, D. Englund, M. Soljačić, C. Denz, D. A. B. Miller, and D. Psaltis, Inference in artificial intelligence with deep optics and photonics, *Nature* **588**, 39 (2020).
 - [8] X. Xu, M. Tan, B. Corcoran, J. Wu, A. Boes, T. G. Nguyen, S. T. Chu, B. E. Little, D. G. Hicks, R. Morandotti, A. Mitchell, and D. J. Moss, 11 tops photonic convolutional accelerator for optical neural networks, *Nature* **589**, 44 (2021).
 - [9] J. Feldmann, N. Youngblood, M. Karpov, H. Gehring, X. Li, M. Stappers, M. Le Gallo, X. Fu, A. Lukashchuk, A. S. Raja, J. Liu, C. D. Wright, A. Sebastian, T. J. Kippenberg, W. H. P. Pernice, and H. Bhaskaran, Parallel convolutional processing using an integrated photonic tensor core, *Nature* **589**, 52 (2021).
 - [10] T. Zhou, X. Lin, J. Wu, Y. Chen, H. Xie, Y. Li, J. Fan, H. Wu, L. Fang, and Q. Dai, Large-scale neuromorphic optoelectronic computing with a reconfigurable diffractive processing unit, *Nature Photonics* **15**, 367 (2021).
 - [11] T. Wang, S.-Y. Ma, L. G. Wright, T. Onodera, B. C. Richard, and P. L. McMahon, An optical neural network using less than 1 photon per multiplication, *Nature Communications* **13**, 123 (2022).
 - [12] F. Ashtiani, A. J. Geers, and F. Aflatouni, An on-chip photonic deep neural network for image classification, *Nature* **606**, 501 (2022).
 - [13] A. Sludds, S. Bandyopadhyay, Z. Chen, Z. Zhong, J. Cochrane, L. Bernstein, D. Bunandar, P. B. Dixon, S. A. Hamilton, M. Streshinsky, A. Novack, T. Baehr-Jones, M. Hochberg, M. Ghobadi, R. Hamerly, and D. Englund, Delocalized photonic deep learning on the internet's edge, *Science* **378**, 270 (2022).
 - [14] B. J. Shastri, A. N. Tait, T. Ferreira de Lima, W. H. P. Pernice, H. Bhaskaran, C. D. Wright, and P. R. Prucnal, Photonics for artificial intelligence and neuromorphic computing, *Nature Photonics* **15**, 102 (2021).
 - [15] R. Hamerly, L. Bernstein, A. Sludds, M. Soljačić, and D. Englund, Large-scale optical neural networks based on photoelectric multiplication, *Phys. Rev. X* **9**, 021032 (2019).
 - [16] M. A. Nahmias, T. F. de Lima, A. N. Tait, H.-T. Peng, B. J. Shastri, and P. R. Prucnal, Photonic multiply-accumulate operations for neural networks, *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1 (2020).
 - [17] L. Bernstein, A. Sludds, C. Panuski, S. Trajtenberg-Mills, R. Hamerly, and D. Englund, Single-shot optical neural network, *Science Advances* **9**, eadg7904 (2023).
 - [18] A. Varri, F. Brücknerhoff-Plückelmann, J. Dijkstra, D. Wendland, R. Bankwitz, A. Agnihotri, and W. H. P. Pernice, Noise-resilient photonic analog neural networks, *J. Lightwave Technol.* **42**, 7969 (2024).

- [19] A. Lugnan, S. Aggarwal, F. Brücknerhoff-Plückelmann, C. D. Wright, W. H. P. Pernice, H. Bhaskaran, and P. Bienstman, Emergent self-adaptation in an integrated photonic neural network for backpropagation-free learning, *Advanced Science* **12**, 2404920 (2025).
- [20] S. Choi, Y. Salamin, C. Roques-Carnes, R. Dangovski, D. Luo, Z. Chen, M. Horodyski, J. Sloan, S. Z. Uddin, and M. Soljačić, Photonic probabilistic machine learning using quantum vacuum noise, *Nature Communications* **15**, 7760 (2024).
- [21] S.-Y. Ma, T. Wang, J. Laydevant, L. G. Wright, and P. L. McMahon, Quantum-limited stochastic optical neural networks operating at a few quanta per activation, *Nature Communications* **16**, 359 (2025).
- [22] A. N. Tait, T. Ferreira de Lima, M. A. Nahmias, H. B. Miller, H.-T. Peng, B. J. Shastri, and P. R. Prucnal, Silicon photonic modulator neuron, *Phys. Rev. Appl.* **11**, 064043 (2019).
- [23] P. Stark, F. Horst, R. Dangel, J. Weiss, and B. J. Offrein, Opportunities for integrated photonic neural networks, *Nanophotonics* **9**, 4221 (2020).
- [24] W. Bogaerts, D. Pérez, J. Capmany, D. A. B. Miller, J. Poon, D. Englund, F. Morichetti, and A. Melloni, Programmable photonic circuits, *Nature* **586**, 207 (2020).
- [25] C. Wu, H. Yu, S. Lee, R. Peng, I. Takeuchi, and M. Li, Programmable phase-change metasurfaces on waveguides for multimode photonic convolutional neural network, *Nature Communications* **12**, 96 (2021).
- [26] M. Miscuglio, Z. Hu, S. Li, J. K. George, R. Capanna, H. Dalir, P. M. Bardet, P. Gupta, and V. J. Sorger, Massively parallel amplitude-only fourier neural network, *Optica* **7**, 1812 (2020).
- [27] J. W. Goodman, A. R. Dias, and L. M. Woody, Fully parallel, high-speed incoherent optical method for performing discrete fourier transforms, *Opt. Lett.* **2**, 1 (1978).
- [28] D. Psaltis, D. Brady, and K. Wagner, Adaptive optical networks using photorefractive crystals, *Appl. Opt.* **27**, 1752 (1988).
- [29] J. Dong, M. Rafayelyan, F. Krzakala, and S. Gigan, Optical reservoir computing using multiple light scattering for chaotic systems prediction, *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1 (2020).
- [30] J. Chang, V. Sitzmann, X. Dun, W. Heidrich, and G. Wetzstein, Hybrid optical-electronic convolutional neural networks with optimized diffractive optics for image classification, *Scientific Reports* **8**, 12324 (2018).
- [31] M. W. Matthès, P. del Hougne, J. de Rosny, G. Lerosey, and S. M. Popoff, Optical complex media as universal reconfigurable linear operators, *Optica* **6**, 465 (2019).
- [32] J. Bueno, S. Maktoobi, L. Froehly, I. Fischer, M. Jacquot, L. Larger, and D. Brunner, Reinforcement learning in a large-scale photonic recurrent neural network, *Optica* **5**, 756 (2018).
- [33] J. Spall, X. Guo, T. D. Barrett, and A. I. Lvovsky, Fully reconfigurable coherent optical vector-matrix multiplication, *Opt. Lett.* **45**, 5752 (2020).
- [34] Y. Liao, M.-H. Hsieh, and C. Ferrie, Quantum optimization for training quantum neural networks, *Quantum Machine Intelligence* **6**, 33 (2024).
- [35] J. Wen, Z. Huang, D. Cai, and L. Qian, Enhancing the expressivity of quantum neural networks with residual connections, *Communications Physics* **7**, 220 (2024).
- [36] L. Bischof, S. Teodoropol, R. M. Füchslin, and K. Stockinger, Hybrid quantum neural networks show strongly reduced need for free parameters in entity matching, *Scientific Reports* **15**, 4318 (2025).
- [37] H. Hirai, Practical application of quantum neural network to materials informatics, *Scientific Reports* **14**, 8583 (2024).
- [38] J. Thompson, Quantum neural networks can be normal, *Nature Physics* **21**, 1042 (2025).
- [39] D. García-Martín, M. Larocca, and M. Cerezo, Quantum neural networks form gaussian processes, *Nature Physics* **21**, 1153 (2025).
- [40] A. Chen and M. Heyl, Empowering deep neural quantum states through efficient optimization, *Nature Physics* **20**, 1476 (2024).
- [41] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nature Physics* **15**, 1273 (2019).
- [42] M. S. Moreira, G. G. Guerreschi, W. Vlothuizen, J. F. Marques, J. van Straten, S. P. Premaratne, X. Zou, H. Ali, N. Muthusubramanian, C. Zachariadis, J. van Someren, M. Beekman, N. Haider, A. Bruno, C. G. Almudever, A. Y. Matsuura, and L. DiCarlo, Realization of a quantum neural network using repeat-until-success circuits in a superconducting quantum processor, *npj Quantum Information* **9**, 118 (2023).
- [43] B. Zhang, J. Liu, X.-C. Wu, L. Jiang, and Q. Zhuang, Dynamical transition in controllable quantum neural networks with large depth, *Nature Communications* **15**, 9354 (2024).
- [44] X. Pan, Z. Lu, W. Wang, Z. Hua, Y. Xu, W. Li, W. Cai, X. Li, H. Wang, Y.-P. Song, C.-L. Zou, D.-L. Deng, and L. Sun, Deep quantum neural networks on a superconducting processor, *Nature Communications* **14**, 4006 (2023).
- [45] Y. Zhu, Y.-D. Wu, G. Bai, D.-S. Wang, Y. Wang, and G. Chiribella, Flexible learning of quantum states with generative query neural networks, *Nature Communications* **13**, 6222 (2022).
- [46] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf, Training deep quantum neural networks, *Nature Communications* **11**, 808 (2020).
- [47] N. Killoran, T. R. Bromley, J. M. Arrazola, M. Schuld, N. Quesada, and S. Lloyd, Continuous-variable quantum neural networks, *Phys. Rev. Res.* **1**, 033063 (2019).
- [48] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New Journal of Physics* **18**, 023023 (2016).
- [49] T. van der Sar, Z. H. Wang, M. S. Blok, H. Bernien, T. H. Taminiau, D. M. Toyli, D. A. Lidar, D. D. Awschalom, R. Hanson, and V. V. Dobrovitski, Decoherence-protected quantum gates for a hybrid solid-state spin register, *Nature* **484**, 82 (2012).
- [50] X. Xue, M. Russ, N. Samkharadze, B. Undseth, A. Sammak, G. Scappucci, and L. M. K. Vandersypen, Quantum logic with spin qubits crossing the surface code threshold, *Nature* **601**, 343 (2022).
- [51] M. Mićuda, R. Stárek, J. Fiurášek, and R. Filip, Decoherence-resilient linear optical two-qubit quantum gate, *Phys. Rev. Appl.* **14**, 054066 (2020).
- [52] P. Huang, N. M. Zimmerman, and G. W. Bryant, Spin decoherence in a two-qubit phase gate: the critical role of tunneling noise, *npj Quantum Information* **4**, 62 (2018).
- [53] A. Kitaev, Anyons in an exactly solved model and beyond, *Annals of Physics* **321**, 2 (2006), january Special

Issue.

- [54] S. Choi, J. Kim, J. Kwak, N. Park, and S. Yu, Topologically protected all-optical memory, *Advanced Electronic Materials* **8**, 2200579 (2022).
- [55] J. Roffe, Quantum error correction: an introductory guide, *Contemporary Physics* **60**, 226 (2019).
- [56] G. R. Steinbrecher, J. P. Olson, D. Englund, and J. Carolan, Quantum optical neural networks, *npj Quantum Information* **5**, 60 (2019).
- [57] J. Ewaniuk, J. Carolan, B. J. Shastri, and N. Rotenberg, Imperfect quantum photonic neural networks, *Advanced Quantum Technologies* **6**, 2200125 (2023).
- [58] S. Roncallo, A. R. Morgillo, S. Lloyd, C. Macchiavello, and L. Maccone, Quantum optical shallow networks (2025), [arXiv:2507.21036 \[quant-ph\]](#).
- [59] N. Killoran, T. R. Bromley, J. M. Arrazola, M. Schuld, N. Quesada, and S. Lloyd, Continuous-variable quantum neural networks, *Phys. Rev. Res.* **1**, 033063 (2019).
- [60] M. Lewenstein and M. Olko, 'quantum' neural networks, *Network: Computation in Neural Systems* **2**, 207 (1991).
- [61] M. Lewenstein and M. Olko, Storage capacity of "quantum" neural networks, *Phys. Rev. A* **45**, 8938 (1992).
- [62] E. Behrman, J. Steck, and S. Skinner, A spatial quantum neural computer, in *IJCNN'99. International Joint Conference on Neural Networks. Proceedings (Cat. No. 99CH36339)*, Vol. 2 (1999) pp. 874–877 vol.2.
- [63] A. Jennings, P. Horan, and J. Hegarty, Optical neural network with quantum well devices, *Appl. Opt.* **33**, 1469 (1994).
- [64] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, Experimental realization of any discrete unitary operator, *Phys. Rev. Lett.* **73**, 58 (1994).
- [65] N. Matsuda, R. Shimizu, Y. Mitsumori, H. Kosaka, and K. Edamatsu, Observation of optical-fibre kerr nonlinearity at the single-photon level, *Nature Photonics* **3**, 95 (2009).
- [66] T. A. Johnson, E. Urban, T. Henage, L. Isenhower, D. D. Yavuz, T. G. Walker, and M. Saffman, Rabi oscillations between ground and rydberg states with dipole-dipole atomic interactions, *Phys. Rev. Lett.* **100**, 113003 (2008).
- [67] S. Choi, Y. Salamin, C. Roques-Carmes, J. Sloan, M. Horodyski, and M. Soljačić, Observing the dynamics of quantum states generated inside nonlinear optical cavities, *Nature Communications* **16**, 7576 (2025).
- [68] Y. Lu, Y. Huang, J. Cheng, R. Ma, X. Xu, Y. Zang, Q. Wu, and J. Xu, Nonlinear optical physics at terahertz frequency, *Nanophotonics* **13**, 3279 (2024).
- [69] G. Yumoto, R. Matsunaga, H. Hibino, and R. Shimano, Ultrafast terahertz nonlinear optics of landau level transitions in a monolayer graphene, *Phys. Rev. Lett.* **120**, 107401 (2018).
- [70] A. Leitenstorfer, K. A. Nelson, K. Reimann, and K. Tanaka, Focus on nonlinear terahertz studies, *New Journal of Physics* **16**, 045016 (2014).
- [71] A. Stockklauser, P. Scarlino, J. V. Koski, S. Gasparinetti, C. K. Andersen, C. Reichl, W. Wegscheider, T. Ihn, K. Ensslin, and A. Wallraff, Strong coupling cavity qed with gate-defined double quantum dots enabled by a high impedance resonator, *Phys. Rev. X* **7**, 011030 (2017).
- [72] D. Englund, A. Faraon, I. Fushman, N. Stoltz, P. Petroff, and J. Vučković, Controlling cavity reflectivity with a single quantum dot, *Nature* **450**, 857 (2007).
- [73] M. Winger, A. Badolato, K. J. Hennessy, E. L. Hu, and A. m. c. Imamoglu, Quantum dot spectroscopy using cavity quantum electrodynamics, *Phys. Rev. Lett.* **101**, 226808 (2008).
- [74] M. Kornjaca, H.-Y. Hu, C. Zhao, J. Wurtz, P. Weinberg, M. Hamdan, A. Zhdanov, S. H. Cantu, H. Zhou, R. A. Bravo, K. Bagnall, J. I. Basham, J. Campo, A. Choukri, R. DeAngelo, P. Frederick, D. Haines, J. Hammett, N. Hsu, M.-G. Hu, F. Huber, P. N. Jepsen, N. Jia, T. Karolyshyn, M. Kwon, J. Long, J. Lopatin, A. Lukin, T. Macri, O. Markovic, L. A. Martinez-Martinez, X. Meng, E. Ostroumov, D. Paquette, J. Robinson, P. S. Rodriguez, A. Singh, N. Sinha, H. Thoreen, N. Wan, D. Waxman-Lenz, T. Wong, K.-H. Wu, P. L. S. Lopes, Y. Boger, N. Gemelke, T. Kitagawa, A. Keesling, X. Gao, A. Bylinskii, S. F. Yelin, F. Liu, and S.-T. Wang, Large-scale quantum reservoir learning with an analog quantum computer (2024), [arXiv:2407.02553 \[quant-ph\]](#).
- [75] C. Zhu, P. J. Ehlers, H. I. Nurdin, and D. Soh, Practical few-atom quantum reservoir computing, *Phys. Rev. Res.* **7**, 023290 (2025).
- [76] C. Zhu, P. J. Ehlers, H. I. Nurdin, and D. Soh, Minimalistic and scalable quantum reservoir computing enhanced with feedback (2024), [arXiv:2412.17817 \[quant-ph\]](#).
- [77] L. Deng, The mnist database of handwritten digit images for machine learning research [best of the web], *IEEE Signal Processing Magazine* **29**, 141 (2012).
- [78] S. Basu, S. Ganguly, S. Mukhopadhyay, R. DiBiano, M. Karki, and R. Nemani, Deepsat: a learning framework for satellite imagery, in *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems*, SIGSPATIAL 15 (Association for Computing Machinery, New York, NY, USA, 2015).
- [79] D. Soh, E. Chatterjee, and M. Eichenfield, High-fidelity state transfer between leaky quantum memories, *Phys. Rev. Res.* **3**, 033027 (2021).
- [80] J. C. Taylor, E. Chatterjee, W. F. Kindel, D. Soh, and M. Eichenfield, Reconfigurable quantum phononic circuits via piezo-acoustomechanical interactions, *npj Quantum Information* **8**, 19 (2022).
- [81] A. E. Bazarov, I. S. Goldobin, P. G. Eliseev, O. A. Kobilzhanov, G. T. Pak, T. V. Petrakova, T. N. Pushkina, and A. T. Semenov, Phase locking of stimulated emission from arrays of stripe gaaias/gaas lasers using active directional couplers, *Soviet Journal of Quantum Electronics* **17**, 551 (1987).
- [82] V. V. Likhanskiĭ, Stability of phase-locked stimulated emission from optically coupled lasers with a nonlinear active medium, *Soviet Journal of Quantum Electronics* **21**, 613 (1991).
- [83] N. N. Ledentsov, V. A. Shchukin, M. Grundmann, N. Kirstaedter, J. Böhrer, O. Schmidt, D. Bimberg, V. M. Ustinov, A. Y. Egorov, A. E. Zhukov, P. S. Kop'ev, S. V. Zaitsev, N. Y. Gordeev, Z. I. Alferov, A. I. Borovkov, A. O. Kosogov, S. S. Ruvimov, P. Werner, U. Gösele, and J. Heydenreich, Direct formation of vertically coupled quantum dots in stranski-krastanow growth, *Phys. Rev. B* **54**, 8743 (1996).
- [84] A. Fedrizzi, R. Ursin, T. Herbst, M. Nespoli, R. Prevedel, T. Scheidl, F. Tiefenbacher, T. Jennewein, and A. Zeilinger, High-fidelity transmission of entanglement over a high-loss free-space channel, *Nature Physics* **5**, 389 (2009).

- [85] D. Bonneau, G. J. Mendoza, J. L. O'Brien, and M. G. Thompson, Effect of loss on multiplexed single-photon sources, [New Journal of Physics](#) **17**, 043057 (2015).
- [86] A. Mao, M. Mohri, and Y. Zhong, Cross-entropy loss functions: theoretical analysis and applications, in *Proceedings of the 40th International Conference on Machine Learning*, ICML'23 (JMLR.org, 2023).
- [87] H. Yao, D.-l. Zhu, B. Jiang, and P. Yu, Negative log likelihood ratio loss for deep neural network classification, in *Proceedings of the Future Technologies Conference (FTC) 2019*, edited by K. Arai, R. Bhatia, and S. Kapoor (Springer International Publishing, Cham, 2020) pp. 276–282.