Quantum Machine Learning on the Photonic Platform

Denis Stanev[1,*], Nicolò Spagnolo[2] and Fabio Sciarrino[2]

[1]Gran Sasso Science Institute, Viale Francesco Crispi 7, I-67100 L'Aquila, Italy
[2]Dipartimento di Fisica, Sapienza Università di Roma, Piazzale Aldo Moro 5, I-00185 Roma, Italy

Abstract

Here, we present two of our works focusing on Quantum Machine Learning. We discuss both about the capabilities of trained quantum optical chips, as well as about methods to train them efficiently. We also consider how the training process will be affected by limitations and imperfections in a possible physical implementation, and show how the performance of the trained circuit is impacted by these issues, as well as how the training process can partially overcome these additional challenges. While these works focus on the quantum optical platform, many of the aspects that we discussed in these works also apply to other quantum computing technologies.

Keywords

Quantum Machine Learning, Quantum Optics, Quantum Computing, Machine Learning

1. Introduction

Quantum devices and quantum algorithms promise substantial advantages in many computational tasks, with even a theoretical exponential speedup when compared to classical computers on specific tasks [1, 2, 3]. In recent years, significant advancements have been made in this field, with the first experiments [4, 5] tackling the quantum advantage regime. At the same time, machine learning, and more specifically deep learning [6], has been finding use in many different sectors. Thus, it is no surprise that the meeting point of these 2 fields, Quantum Machine Learning, has become a topic of interest.

Here, we present some of our work pertaining to Quantum Machine Learning (QML) over the last couple of years. The main focus will be on 2 of our projects: the first one is about training a Quantum Optical Neural Network (QONN) to perform deterministic optimal quantum cloning [7] (as opposed to the previously known optical quantum cloners, which are probabilistic [8, 9]), while the second project addresses the training of a continuously coupled optical chip using machine learning techniques to perform a desired unitary transformation.

Both these works relate to training quantum optical chips for specific tasks, with the first one being more focused on the capabilities of a trained chip, while the second one being more focused on how a quantum optical chip can be efficiently trained. It is also worth noting that, while these works focus on the quantum optical platform, many of the aspects that we discussed in these works also apply to other quantum computing technologies.

2. Training a deterministic universal optical quantum cloner via QONNs

The QONN is the quantum optical extension of neural networks, and is composed of multiple layers, each composed of a linear unitary transformation and a set of nonlinear single-site transformations. Input states in the network are photon states employed to encode qubits. More specifically, the dual-rail encoding is used to associate each pair of modes to a qubit, where states $|0\rangle$ and $|1\rangle$ correspond to a single photon in the first and second mode respectively. The QONN can be considered the quantum optical version of the multi-layer perceptron (MLP), and was first introduced in a work by Steinbrecher et al. [10].

AIQxQIA 2025: International Workshop on AI for Quantum and Quantum for AI | co-located with ECAI 2025, Bologna, Italy
*Corresponding author.
denis.stanev@gssi.it (D. Stanev)

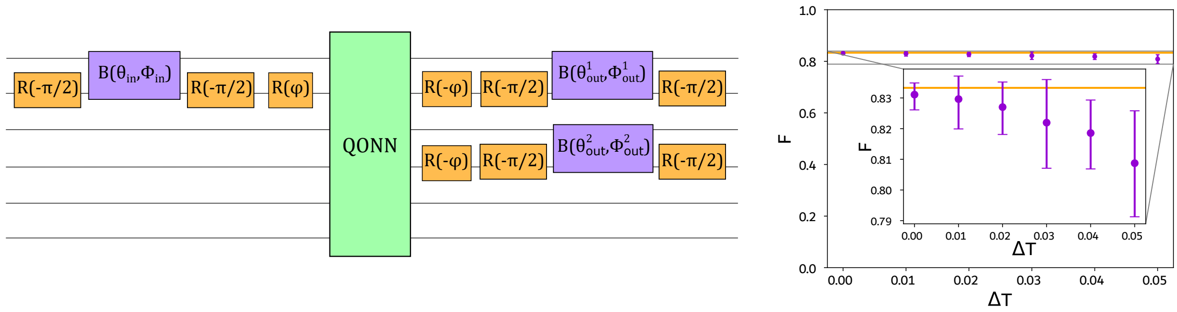

Figure 1: Left, complete circuit including the input state preparation stage and the output state measurement one on each pair of modes corresponding to a logical qubit. **Right**, fidelity of the clones as a function of the fabrication error ΔT_j in the directional couplers. The purple bars represent the range between the minimum and maximum fidelity for each QONN, while the orange line represents the theoretical maximum average fidelity.

In this work our objective was to train a simulated QONN to obtain a deterministic optimal quantum cloner for the optical platform. We managed to successfully do so, and obtained an optical quantum cloner whose performance is very close to that of an optimal universal quantum cloner. The optical quantum cloner obtained was deterministic, which is very important, as current optical implementations are inherently probabilistic [8, 9]. We also verified numerically the robustness of this approach with respect to experimental imperfections, mainly relating to the effects of a restricted range of transmittivity for the tunable beam-splitters on the training, which is relevant for potential future implementations. In Figure 1 we show the structure of the circuit we used to train the QONN, as well as the results of the fidelity of the trained QONNs in the case of simulated experimental imperfections of various magnitude.

This work thus showed the capabilities and flexibility of QONNs, which might prove important for designing even more complex optical circuits, as well as highlighting some potential uses for QONNs.

3. Training a continuously coupled optical chip via machine learning techniques

Continuously coupled optical chips [11, 12] offer several potential advantages over other options, such as shorter circuits, and thus lower losses, due to their more compact nature, especially when considering 3D designs [11]. However, one issue that this type of quantum optical chip currently presents is that, given a specific tunable continuously coupled optical chip, it is not yet known how to obtain the parameters that achieve a specific unitary transformation, while in the case of quantum optical chips with discrete points of interaction such decompositions are known [13, 14].

Thus, in this work we make use of machine learning techniques to train the parameters of a continuously coupled chip to perform a desired unitary transformation. We simulate different sizes and structures of continuously coupled optical chips, and check how these techniques work in these various scenarios. We consider both flat and 3D designs, as well as a varying number of optical modes and the effect on training of having limited control on the quantum optical chip. We also evaluate the errors caused by several potential experimental limitations, such as those produced by a limited number of samples, and those caused by imperfections in the parametrization of the chip. We also tackle the issue of training the quantum optical chip with constraints similar as those that would be present in an actual physical experiment, such as the fact that backpropagation cannot be used [15], and thus alternative strategies need to be employed to obtain an estimation of the gradient. On top of that, we implement various optimizations during the training process that would reduce the number of measurements needed to physically implement this training process, thus leading to faster training times. We also train a classical machine learning model, to work in tandem with the quantum machine learning process, in cases in which only training the quantum chip directly does not provide the desired results.

This work thus showed the challenges that arise when attempting to train a quantum chip in a realistic scenario, as well as giving some insights on methods and optimizations that can be used to improve and accelerate the training process.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, SIAM Journal on Computing 26 (1997) 1484–1509. URL: https://doi.org/10.1137/S0097539795293172. doi:10.1137/S0097539795293172. arXiv:https://doi.org/10.1137/S0097539795293172.

[2] R. Babbush, D. W. Berry, R. Kothari, R. D. Somma, N. Wiebe, Exponential quantum speedup in simulating coupled classical oscillators, Physical Review X 13 (2023). URL: http://dx.doi.org/10.1103/PhysRevX.13.041041. doi:10.1103/physrevx.13.041041.

[3] N. Spagnolo, D. J. Brod, E. F. Galvão, F. Sciarrino, Non-linear boson sampling, npj Quantum Information 9 (2023) 3. doi:10.1038/s41534-023-00676-x.

[4] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffman, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Motus, O. Naaman, M. Neeley, C. Neill, M. Yuezhen Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, J. M. Martinis, Quantum supremacy using a programmable superconducting processor, Nature 574 (2019) 505–510. URL: http://dx.doi.org/10.1038/s41586-019-1666-5. doi:10.1038/s41586-019-1666-5.

[5] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, J.-W. Pan, Quantum computational advantage using photons, Science 370 (2020) 1460–1463. doi:10.1126/science.abe8770.

[6] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature 521 (2015) 436–444. URL: https://doi.org/10.1038/nature14539. doi:10.1038/nature14539.

[7] D. Stanev, N. Spagnolo, F. Sciarrino, Deterministic optimal quantum cloning via a quantum-optical neural network, Phys. Rev. Res. 5 (2023) 013139. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.5.013139. doi:10.1103/PhysRevResearch.5.013139.

[8] L. ye, X.-k. Song, J. Yang, Q. Yang, Y.-C. Ma, Optical implementation of the optimal universal and phase-covariant quantum cloning machines, Modern Physics Letters B 26 (2012) 50102–. doi:10.1142/S0217984912501023.

[9] F. D. Martini, D. Pelliccia, F. Sciarrino, Contextual, optimal, and universal realization of the quantum cloning machine and of the not gate, Phys. Rev. Lett. 92 (2004) 067901. URL: https://link.aps.org/doi/10.1103/PhysRevLett.92.067901. doi:10.1103/PhysRevLett.92.067901.

[10] G. R. Steinbrecher, J. P. Olson, D. Englund, J. Carolan, Quantum optical neural networks, npj Quantum Inf. 5 (2019) 60. doi:10.1038/s41534-019-0174-7.

[11] F. Hoch, S. Piacentini, T. Giordani, Z.-N. Tian, M. Iuliano, C. Esposito, A. Camillini, G. Carvacho, F. Ceccarelli, N. Spagnolo, A. Crespi, F. Sciarrino, R. Osellame, Reconfigurable continuously-coupled 3d photonic circuit for boson sampling experiments, npj Quantum Information 8 (2022). URL: http://dx.doi.org/10.1038/s41534-022-00568-6. doi:10.1038/s41534-022-00568-6.

[12] Y. Yang, R. J. Chapman, B. Haylock, F. Lenzini, Y. N. Joglekar, M. Lobino, A. Peruzzo, Programmable high-dimensional hamiltonian in a photonic waveguide array, Nature Communications 15 (2024). URL: http://dx.doi.org/10.1038/s41467-023-44185-z. doi:10.1038/s41467-023-44185-z.

[13] M. Reck, A. Zeilinger, H. J. Bernstein, P. Bertani, Experimental realization of any discrete unitary operator, Phys. Rev. Lett. 73 (1994) 58–61. URL: https://link.aps.org/doi/10.1103/PhysRevLett.73.58. doi:10.1103/PhysRevLett.73.58.

[14] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, I. A. Walmsley, Optimal design for universal multiport interferometers, Optica 3 (2016) 1460–1465. URL: https://opg.optica.org/optica/abstract.cfm?URI=optica-3-12-1460. doi:10.1364/OPTICA.3.001460.

[15] A. Abbas, R. King, H.-Y. Huang, W. J. Huggins, R. Movassagh, D. Gilboa, J. R. McClean, On quantum backpropagation, information reuse, and cheating measurement collapse, 2023. URL: https://arxiv.org/abs/2305.13362. arXiv:2305.13362.