Photonic Quantum-Accelerated Machine Learning

Markus Rambach[1], Abhishek Roy[1,2], Alexei Gilchrist[2], Akitada Sakurai[3],

William J. Munro[3], Kae Nemoto[3], and Andrew G. White[1]

[1] *School of Mathematics and Physics, The University of Queensland, Brisbane, QLD 4072, Australia*

[2] *Department of Physics and Astronomy, Macquarie University, Sydney, NSW 2113, Australia*

[3] *Okinawa Institute of Science and Technology Graduate University, Onna-son, Okinawa 904-0495, Japan*

Machine learning (ML) and quantum physics are two of the most transformative fields in the modern era of science and technology. ML is widely applied in our society, but has yet to broadly leverage the unique benefits of quantum resources. Boson sampling—a quantum interference-based sampling protocol—is a resource that is classically hard to simulate and can be implemented on current photonic hardware. Here, we introduce quantum optical reservoir computing (QORC) [1], a photonic accelerator for classical ML, in which a boson-sampling interferometer acts as a fixed, high-dimensional nonlinear quantum reservoir, enabling quantum-enhanced feature transformation with minimal training overhead.

In QORC, input images are first compressed using principal component analysis. The first M components—matching the number of utilised modes—are then encoded into phases of an integrated photonic circuit. Finally, an N-photon input is processed by the chip via multi-photon interference and measured by bucket detectors. The resulting photon patterns form quantum fingerprints that are fed to a simple classical linear classifier. The interferometer remains fully untrained, preserving the simplicity and scalability of reservoir computing while providing strong nonlinear separability.

We benchmark QORC on handwritten-digit and biomedical image datasets, comparing our scheme to classical baselines and small shallow neural networks. We numerically study configurations with N=1-5 photons in M=12/20/24 modes, and experimentally implement N=3, M=12 on a real-world photonic quantum processing unit. We demonstrate consistent advantages across data efficiency, accuracy, robustness to photon imperfections (e.g. indistinguishability), severe class imbalances and sparse data, and real-device performance. The improvement in data efficiency is especially noteworthy: QORC achieves classical linear-classifier performance while using roughly twenty times less data in training. Crucially, we demonstrate the acceleration and scalability of our scheme on photonic quantum hardware, providing experimental validation that boson-sampling-enhanced learning yields real performance gains on actual devices [2].

These results demonstrate that quantum photonic processors can accelerate practical ML workloads today. Scaling photonic platforms, improved sampling, and applying QORC to temporal tasks represent natural directions for extending quantum reservoir computing.

References

[1] A. Sakurai, A. Hayashi, W. J. Munro, and K. Nemoto, Optica Quantum 3(3), 238 (2025).

[2] M. Rambach, A. Roy, A. Gilchrist, A. Sakurai, W. J. Munro, K. Nemoto, and A. G. White, eprint arXiv:2512.08318 (2025).