Abstract

This dissertation investigates hardware-aware Quantum Machine Learning under Noisy Intermediate-Scale Quantum constraints, addressing the gap between idealized algorithm design and physical hardware limitations, including decoherence, gate infidelity, and restricted connectivity. The central claim is that practical progress depends on co-design between high-level models and low-level control.

First, the dissertation establishes the theoretical foundations of encoding, kernels, and trainability, then applies them to a neutral-atom Quantum Kernel Estimation pipeline. On the benchmark used, the proposed quantum kernel reaches a mean accuracy of $79.8\% \pm 8.8\%$ across 30 independent dataset seeds, significantly outperforming the classical radial-basis-function baseline ($52.8\% \pm 10.2\%$) with a mean accuracy gap of $27.0\% \pm 13.4\%$ (paired t-test p-value of 6.33×10^{-12}), while highlighting the role of geometric connectivity, pulse-level gate synthesis, and hardware clock quantization.

Second, for superconducting qubit control, the dissertation introduces Adapted Randomized Benchmarking (ARB) as a statistically rigorous method for both estimating gate-set fidelity and serving as a loss function. A compact neural network is first pre-trained with a standard MSE loss on simulation data to produce high-fidelity single-qubit pulse coefficients. The physical parameters of the simulated system are then deliberately altered, emulating realistic drifts that occur after deployment, such as changes in anharmonicity. Rather than retraining from scratch, the network is fine-tuned on the shifted parameters using the ARB fidelity estimate as the loss function. The fine-tuning procedure restores the fidelity in a statistically rigorous manner, demonstrating that lightweight on-device recalibration is feasible and confirming the viability of embedded real-time pulse calibration on resource-constrained platforms such as FPGAs.

Third, in the photonic domain, a gate-based quantum circuit is shown to reproduce with high fidelity the dynamics of a single-photon protective-measurement experiment. The simulation faithfully captures the Quantum-Zeno-inspired survival-probability enhancement observed in the laboratory, with the protective configuration reducing the percentage relative error by approximately 83% with respect to the non-protective one despite increased circuit depth. This validates the gate-model decomposition of the spatial translation and provides a tool for designing protocol variants or porting the scheme to platforms lacking native continuous-variable operations.

Overall, the results support a unified roadmap in which platform-specific control, model design, and measurement strategy jointly improve robustness and prepare the ground for future large-scale QML validation.