

Deployable Quantum Machine Learning by Design: Hardware-Aware and Adversarially Robust Neural Architecture Search for Quantum Circuits

Walid El Maouaki^{1,2}, Muhammad Kashif^{1,3}, Muhammad Shafique^{1,2,3}

¹ eBrain Lab, Division of Engineering, New York University Abu Dhabi, PO Box 129188, Abu Dhabi, UAE

² Center for Cyber Security, NYUAD Research Institute, New York University Abu Dhabi, UAE

³ Center for Quantum and Topological Systems, NYUAD Research Institute, New York University Abu Dhabi, UAE

Abstract: Quantum machine learning (QML) models are often optimized primarily for accuracy. However, high performance in terms of accuracy alone does not guarantee practical deployment on near-term quantum hardware. In real execution settings, a quantum circuit must undergo transpilation to the specific quantum backend, which transforms the quantum circuit such that it respects the hardware constraints of the target devices such as qubit connectivity and native gate set. This process often results in increased number of hardware resources and noise accumulation. Additionally, maintaining trustworthy behavior under adversarial or distribution-shifted inputs is of significant importance. Hence, all these factors should be jointly considered while designing QML models. This work proposes a deployability-aware neural architecture search framework which automatically identifies quantum circuits optimized in terms of accuracy, adversarial robustness, and hardware resource usage (after transpilation). A large number of candidate architectures are evaluated after transpilation and tested under noise models and adversarial perturbations. Our framework then produces a set of compact and reliable QML models which exhibit best trade-offs in terms of hardware compatibility and resource overhead, accuracy, and stronger robustness in realistic execution settings.

Introduction and Motivation: Many QML models perform well in noiseless simulation but become impractical after transpilation or noisy execution. This creates a gap between theoretical performance and real hardware usability. This work addresses that gap by shifting quantum circuit design from accuracy-only optimization to deployability-centered optimization.

Methods or Approach: We propose a hardware-aware NAS framework that searches over feature maps, ansatz structures, entanglement patterns, circuit depth, qubit layout, and measurement strategies. Each candidate circuit is evaluated using a multi-objective function that combines clean accuracy, adversarial robustness (incorporated via adversarial risk terms and attack-aware training objectives), noisy-simulation performance under backend-calibrated error models, transpiled depth, two-qubit gate count, SWAP overhead, estimated hardware error, shot efficiency, and optional error-mitigation-aware execution. Tools such as Qiskit, Qiskit Aer, PennyLane, Optuna, and adversarial attack methods can be used to implement and evaluate the framework.

Outcomes: The framework produces QML architectures that maintain competitive accuracy while reducing transpiled depth, two-qubit gate overhead, and sensitivity to hardware noise. The selected models should also show improved robustness against adversarial perturbations compared with accuracy-only circuit designs.

Conclusion and Significance: This work reframes QML circuit design as a deployability-aware optimization problem under realistic deployment conditions. By jointly considering accuracy, hardware constraints, noise resilience, and adversarial robustness, the framework positions quantum architecture search as a bridge between quantum machine learning research and quantum-ready applications. As a result, it supports the development of QML models that are more realistic, reliable, and ready for deployment on near-term quantum devices.

Keywords: quantum machine learning, neural architecture search, hardware-aware optimization, quantum transpilation, adversarial robustness, noise resilience.