

Reinforcement-Guided Adaptive Aggregation for Distributed Quantum Machine Learning

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

¹eBRAIN Lab, Division of Engineering, New York University Abu Dhabi (NYUAD), Abu Dhabi, UAE

²Center for Quantum and Topological Systems (CQTS), NYUAD Research Institute, NYUAD, Abu Dhabi, UAE

³Center for Cyber Security, NYUAD Research Institute, New York University Abu Dhabi, UAE

nouhaila.innan@nyu.edu, walid.el.maouaki@nyu.edu, muhammad.shafique@nyu.edu

ABSTRACT

Distributed quantum machine learning (DQML) provides a promising route for training quantum and hybrid quantum-classical models across multiple clients while keeping data local. This setting is especially relevant when data are heterogeneous and distributed across institutions, devices, or quantum processors, and direct data sharing is restricted or inefficient. However, the performance of DQML depends not only on the local quantum models but also on how client updates are aggregated at the server. A fixed aggregation rule, such as simple averaging, may perform well under balanced and clean settings, but it can become less effective when clients have non-IID data, different sample sizes, noisy updates, unequal circuit depths, hardware-dependent errors, or unstable training behavior, or even faulty or adversarial clients. These challenges are even more pronounced in quantum learning, where shot noise, barren-plateau effects, circuit expressivity, and device noise can cause large variation in update quality across clients and communication rounds.

In this work, we propose an adaptive aggregation-selection framework for DQML based on reinforcement learning (RL) and quantum reinforcement learning (QRL). Instead of using the same aggregation strategy throughout training, we formulate aggregation as a sequential decision-making problem. At each communication round, an agent observes the current training state and selects the aggregation technique that is expected to produce the most reliable global update. The state representation can include client-level losses, validation accuracy, update divergence, parameter variation, client participation, gradient statistics, model confidence, and quantum resource indicators such as circuit depth, number of shots, gate counts, and estimated noise levels. The action space can include several aggregation rules, including FedAvg, weighted averaging, performance-aware aggregation, median-based aggregation, trimmed mean, uncertainty-guided aggregation, and reliability-weighted aggregation. The agent is trained to select actions that maximize long-term global model performance and stability, including resilience against adversarial or unreliable clients.

The proposed framework studies both classical RL and QRL agents. Classical RL agents provide a strong and stable baseline for learning aggregation policies under heterogeneous and adversarial client behavior, while QRL agents use parameterized quantum circuits to represent the policy or value function and explore whether quantum models can support adaptive control in distributed learning. This design allows us to separate two questions: whether adaptive aggregation improves DQML compared with fixed aggregation, and whether QRL offers any advantage over classical RL for learning the aggregation policy. By dynamically selecting aggregation rules during training, the framework aims to improve convergence stability, reduce the effect of unreliable client updates, increase robustness under non-IID data, poisoned or malicious updates, and noisy quantum settings, and improve final model performance. The study positions aggregation in DQML as an intelligent control problem and provides a foundation for comparing fixed, classical RL-based, and QRL-based aggregation strategies in quantum federated and distributed learning systems.

INDEX TERMS

Distributed quantum machine learning, quantum federated learning; adaptive aggregation; reinforcement learning; quantum reinforcement learning; aggregation selection; hybrid quantum-classical learning