

Empirical evidence of generalization in quantum machine learning in small-data regime

Leena Anthony,^{1,2} Artemiy Burov,^{1,2} Nicolas Piro,³ and Clément Javerzac^{1,2,3,*}

¹*School of Life Sciences, FHNW, CH-4132 Muttens, Switzerland* ²*EPFL, CH-1015 Lausanne, Switzerland*

³*MatterDecoder, CH-3008 Bern, Switzerland*

Introduction. Generalization, the ability of a model to perform well on unseen data, is a persisting challenge in machine learning. Classical deep networks excel at image analysis but require large training sets for stable generalization [1]; in medical domain, where privacy, annotation cost, and clinical heterogeneity limit data, they often overfit. Quantum machine learning offers a different inductive bias: Caro *et al.* [2] bounded the generalization error of variational quantum models as $\mathcal{O}(\sqrt{T/N})$ for T trainable gates and N samples (logarithmic in gate reuse), a regime naturally realized by quantum convolutional neural networks (QCNNs) [3] through parameter sharing and demonstrated on digits by Kottmann *et al.* [4]. Recent work cautions, however, that such bounds cannot alone explain generalization: quantum models can fit randomized labels [5] and QCNNs are classically simulable on simple benchmarks [6]. We study QCNN generalization and scalability empirically.

Methods. The QCNN alternates convolutional layers (single-qubit U_3 rotations and translationally invariant R_{XX}, R_{YY}, R_{ZZ} sweeps; 15 shared angles in the first layer, 9 thereafter), measurement-based pooling (mid-circuit measurement with conditional rotations, hardware-executable [7]), and a dense $SU(4)$ layer (15 angles). Parameter sharing fixes each layer’s count, so the total T grows only logarithmically with input size: $T = 45$ for the 6-qubit (8×8) model and $T = 57$ for the 10-qubit (28×28) model, the regime in which Caro *et al.* guarantee bounded error from $\mathcal{O}(\log^2 n)$ samples. *Digits* (scikit-learn, 0 vs. 1, 8×8) ran as a 6-qubit dynamic circuit in Qiskit [8] on the FakeKyiv noise-aware simulator (finite-difference gradients, Adam optimizer, 50 epochs, $N \in \{2, 5, 10, 20, 40, 80\}$, a re-drawn 100-sample test set, 10 repetitions); *BreastMNIST* (MedMNIST v2 [9], malignant vs. normal/benign, 28×28) used a 10-qubit amplitude embedding simulated in PennyLane by exact statevector evaluation (analytic gradients, 100 epochs, $N \in \{20, 40, 80, 160, 294\}$, a fixed balanced 84-sample test set, 20 repetitions).

Results. *Digits*: test accuracy exceeds 90% already at $N = 10$ with a modest train - test gap, and the generalization error falls by more than an order of magnitude ($\sim 4 \times 10^{-2} \rightarrow \sim 10^{-3}$), consistent with Caro *et al.* *BreastMNIST*: absolute performance is lower but consistently above chance, with aggregate balanced accuracy rising from 58.6% ($N = 20$) to 61.9% ($N = 294$) and the loss-based generalization gap falling from $\sim 5.2 \times 10^{-2}$ to $\sim 1.7 \times 10^{-2}$; both classes are recalled above chance (malignant 54–60%, normal/benign 61–65%), so the model does not collapse to a single class. *vs. CNN*: at a matched parameter budget the QCNN learns while an equally small CNN stays at chance; the unconstrained CNN ($\sim 25,666$ parameters) is strongest overall, yet on BreastMNIST the 57-parameter QCNN is competitive in the most data-starved regime, exceeding it at $N = 20$. *Scaling*: amplitude encoding stays qubit-efficient (18 qubits at 512×512 , within IBM Condor’s 1121 limit) but reaches transpiled depth $> 10^6$, whereas angle encoding stays shallow but needs one qubit per pixel (262,144 at 512×512).

Discussion and conclusion. For QCNNs, learning from few samples is attainable in practice. These results have motivated our ongoing research in data-scarce regimes such as clinical trials and mRNA codon optimization, where labelled data are costly or intrinsically limited. Consistent with recent cautions [6], we frame them as evidence of data efficiency rather than as a demonstration of quantum advantage.

* clement.javerzac@fhnw.ch *Original, unpublished results; a full preprint is in preparation.*

References

- [1] S.-C. Huang *et al.*, npj Digit. Med. **6**, 74 (2023).
- [2] M. C. Caro *et al.*, Nat. Commun. **13**, 4919 (2022).
- [3] I. Cong, S. Choi, and M. D. Lukin, Nat. Phys. **15**, 1273 (2019).
- [4] K. Kottmann, L. Mantilla Calderón, and M. Weber, “Generalization in QML from Few Training Data,” PennyLane demo (2022).
- [5] E. Gil-Fuster, J. Eisert, and C. Bravo-Prieto, Nat. Commun. **15**, 2277 (2024).
- [6] P. Bermejo *et al.*, PRX Quantum **7**, 020304 (2026).
- [7] A. Carrera Vazquez *et al.*, Nature **636**, 75 (2024).
- [8] Qiskit contributors, “Qiskit: An Open-Source Framework for Quantum Computing” (2023).
- [9] J. Yang *et al.*, Sci. Data **10**, 41 (2023).