Understanding the power and simulability of Quantum Convolutional Neural Networks

Pablo Bermejo,[1] Paolo Braccia,[1] Manuel S. Rudolph,[2] Zoë Holmes,[2] Lukasz Cincio,[1] and M. Cerezo[1]

[1] *Los Alamos National Laboratory, Los Alamos, NM 87545, USA*

[2] *Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne CH-1015, Switzerland*

Introduction: Quantum Convolutional Neural Networks (QCNN) [1] have seen widespread use in Quantum Machine Learning as they posses desirable features such as requiring shallow circuits and not exhibiting barren plateaus [2]. In particular, these models have been extremely successful for classifying phases of matter [3–7]. In this work we argue that most tasks making use of a QCNN can be solved without the need to run a variational circuit on a quantum computer. That is, we provide theoretical and numerical evidence that if a QCNN can solve a given task, then there exists an efficient classical algorithm that can also solve the problem efficiently. These results are based on two key observations: (1) When randomly initialized QCNNs can only "see" their input state locally; and (2) The tasks for which QCNNs have shown heuristic success, allow for a solution (e.g. classification tasks) via local measurements. Our results further strengthen the connection between absence of barren plateaus and classical simulability presented in [8].

QCNN simulability: Consider the following question: *Why do QCNNs not have barren plateaus when randomly initialized?* To understand its answer, let us begin by recalling that a QCNN is composed of alternating layers of convolutional and pooling layers. During a convolutional layer two-qubit gates act on neighboring qubits in a brick-like fashion, while in a pooling layer half of the qubits are traced out. At the end of the QCNN, one measures the expectation value of some local observable. Let $U(\boldsymbol{\theta})$ denote the unitary that implements a QCNN and O the operator that we measure. Then, our first result is as follows:

Result 1 (Informal)**.** *Consider the adjoint action of the QCNN over the measurement, then* $\mathbb{E}_{\boldsymbol{\theta}}[U^\dagger(\boldsymbol{\theta})OU(\boldsymbol{\theta})] \simeq \sum_j c_j P_j$*, where* P_j *is a Pauli of bodyness* $|P_j|$ *(i.e., acting non-trivially on* $|P_j|$ *qubits) and where* $c_j \in \mathcal{O}(1/b^{|P_j|})$ *(for some* $b > 1$*) are coefficients that decrease exponentially with the bodyness of the Pauli.*

Result 1 shows that randomly initialized QCNNs essentially only have support on local operators, i.e., high-bodyness operators have exponentially small contributions and hence do not really see the initial state during the beginning of the training (see Fig. 1). Clearly, this realization on its own does not say much as the optimal classifying measurement $U^\dagger(\boldsymbol{\theta}^*)OU(\boldsymbol{\theta}^*)$ (where $\boldsymbol{\theta}^*$ denotes the optimal set of parameters) could be a global operator, as is expected in most of the cases for phases of matter classification [9]. However, by revisiting several results in the literature, we argue that the following result generally holds.

Result 2 (Informal)**.** *For most classification tasks in the literature there exists sets of optimal parameters* $\boldsymbol{\theta}^*$ *such that* $U^\dagger(\boldsymbol{\theta}^*)OU(\boldsymbol{\theta}^*)$ *has support on local observables, and the model achieves large classification accuracy.*

They key implication of Result 2 is that the QCNN should be classically simulable as we only need to simulate its action on the subspace of local operators. Importantly, such classical simulation still needs a quantum computer as we need to perform classical shadows [10–12] on the quantum data. Critically, these shadows are used to approximate the action of the QCNN on its initial states. More formally, we argue that for tasks where QCNN work, then the algorithms is simulable in CSIM_{QE} [8]. For convenience, we recall that a problem $\mathcal{C}$ is in CSIM_{QE} if a polynomial-time classical algorithm, which can utilize data obtained from quantum devices in an initial data acquisition phase (also in polynomial time), can compute every instance in $\mathcal{C}$.

Numerical results: To showcase our claim, we classically simulate the action of a QCNN via tensor networks for the task of classify between the 2 phases of the bond-alternating XXX Heisenberg model defined by $H = J_1 \sum_{i \text{ even}} S_i S_{i+1} + J_2 \sum_{i \text{ odd}} S_i S_{i+1}$. The ground states are characterized by a trivial phase and a topologically

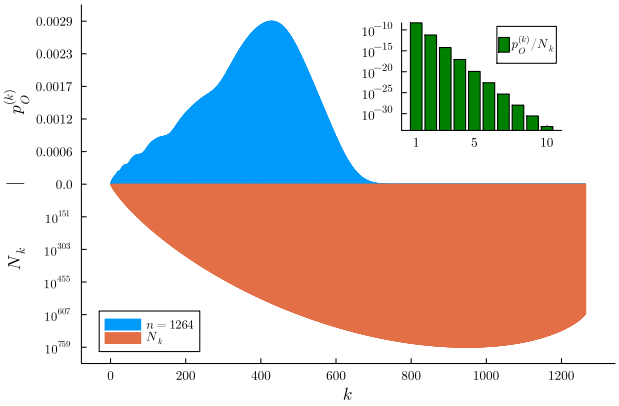

Figure 1: We show the distribution of k-purities ($p_O^{(k)} = \sum_{j\,/\,|P_j|=k} |c_j|^2$) for a QCNN acting on $n = 1264$ qubits (blue), with convolutional layer gates sampled i.i.d from $G = U(4)$. We also present the number of Paulis $N_k = 3^k \binom{n}{k}$ acting on k-qubits (orange). The inset depicts the quotient $p_O^{(k)}/N_k$, which tells us the average weight c_j of a Pauli with bodyness k.

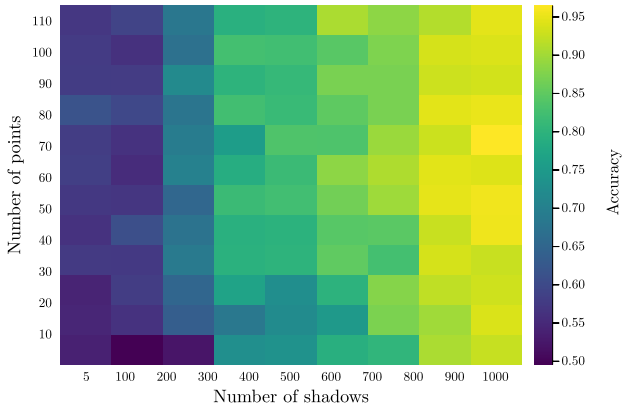

Figure 2: Classification accuracy for a tensor network simulation of a QCNN trained for varying number of data points, and classical shadows taken from each point. We note that the preliminary results shown here corresponds to $n = 6$ qubits, but these will be scaled up for the final poster and for the pre-print.

protected phase (function of the ratio J_1/J_2). We consider a model of 6 qubits (to be extended) and employ a dataset of 500 samples, with equal distribution of both species. Given that the order parameter of this model is local, we expect that a solution can be reached by biasing the tensor network to explore only local operators.

In Fig 2 we show the test accuracy as a function of the 2 parameters defining the computational cost of the data acquisition: number of training points and number of shadows. This plot reveals important features in the training process such as the increase of accuracy as the number of shadows increases, and the slight dependence of the accuracy on the number of training points per phase, proving no overhead in the number of realizations of the quantum circuit using shadow tomography. This trend reflects the expected locality in the classification task involving one trivial phase. The simulated QCNN is able to classify the quantum samples by only making use of local information provided by classical shadows, with no need to rerun an actual quantum circuit for the training stage. The number of shadows needed is comparable to the number of shots employed during ***each optimization step*** in standard training processes, but realized only once in the data acquisition stage. We leave further room for significant performance improvement since fine-tuning a system of 10 qubits provides an accuracy above 95% for just 200 shadows and 40 training points.

Implications and Future Directions: At a fundamental level, our work shows that QCNNs do not need to be run on a quantum device for the tasks where they have shown heuristic success [13–21]. These results, however, are extremely positive from a resource perspective, as training a model based on tomographic data (such as classical shadows) without running on-chip any variational algorithm with expensive optimization loops results in major savings in terms of quantum resources. Moving forward, we intend to show that both standard QCNNs run on a quantum device, as well as their classical-shadows-based surrogates, should fail for classification tasks requiring global operators [22, 23]. To illustrate this limitation, we will consider the problem of distinguishing between two topological phases. This results will ultimately prove that: (1) If the task is easy and the QCNN works, then it can be classically simulated, and replaced by a surrogate; and (2) If the task is hard, then the QCNN will fail and this model should be avoided.

[1] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, Nature Physics **15**, 1273 (2019).

[2] A. Pesah, M. Cerezo, S. Wang, T. Volkoff, A. T. Sornborger, and P. J. Coles, Absence of barren plateaus in quantum convolutional neural networks, Physical Review X **11**, 041011 (2021).

[3] M. C. Caro, H.-Y. Huang, M. Cerezo, K. Sharma, A. Sornborger, L. Cincio, and P. J. Coles, Generalization in quantum machine learning from few training data, Nature Communications **13**, 4919 (2022).

[4] Y.-J. Liu, A. Smith, M. Knap, and F. Pollmann, Model-independent learning of quantum phases of matter with quantum convolutional neural networks, Physical Review Letters **130**, 220603 (2023).

[5] T. Hur, L. Kim, and D. K. Park, Quantum convolutional neural network for classical data classification, Quantum Machine Intelligence **4**, 3 (2022).

[6] J. Herrmann, S. M. Llima, A. Remm, P. Zapletal, N. A. McMahon, C. Scarato, F. Swiadek, C. K. Andersen, C. Hellings, S. Krinner, *et al.*, Realizing quantum convolutional neural networks on a superconducting quantum processor to recognize quantum phases, Nature Communications **13**, 1 (2022).

[7] Q. T. Nguyen, L. Schatzki, P. Braccia, M. Ragone, M. Larocca, F. Sauvage, P. J. Coles, and M. Cerezo, A theory for equivariant quantum neural networks, arXiv preprint arXiv:2210.08566 (2022).

[8] M. Cerezo, M. Larocca, D. García-Martín, N. L. Diaz, P. Braccia, E. Fontana, M. S. Rudolph, P. Bermejo, A. Ijaz, S. Thanasilp, *et al.*, Does provable absence of barren plateaus imply classical simulability? or, why we need to rethink variational quantum computing, arXiv preprint arXiv:2312.09121 (2023).

[9] R. Moessner and J. E. Moore, *Topological Phases of Matter* (Cambridge University Press, 2021).

[10] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nature Physics **16**, 1050 (2020).

[11] H.-Y. Huang, R. Kueng, and J. Preskill, Efficient estimation of pauli observables by derandomization, Physical review letters **127**, 030503 (2021).

[12] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, Provably efficient machine learning for quantum many-body problems, Science **377**, eabk3333 (2022).

[13] A. Letcher, S. Woerner, and C. Zoufal, Tight and efficient gradient bounds for parameterized quantum circuits, arXiv preprint arXiv:2309.12681 (2023).

[14] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, Nature Communications **12**, 1 (2021).

[15] Z. Liu, L.-W. Yu, L.-M. Duan, and D.-L. Deng, The presence and absence of barren plateaus in tensor-network based machine learning, Physical Review Letters **129**, 270501 (2022).

[16] M. S. Rudolph, S. Lerch, S. Thanasilp, O. Kiss, S. Vallecorsa, M. Grossi, and Z. Holmes, Trainability barriers and opportunities in quantum generative modeling, arXiv preprint arXiv:2305.02881 (2023).

[17] L. Schatzki, M. Larocca, Q. T. Nguyen, F. Sauvage, and M. Cerezo, Theoretical guarantees for permutation-equivariant quantum neural networks, npj Quantum Information **10**, 10.1038/s41534-024-00804-1 (2024).

[18] M. T. West, J. Heredge, M. Sevior, and M. Usman, Provably trainable rotationally equivariant quantum machine learning, arXiv preprint arXiv:2311.05873 (2023).

[19] N. L. Diaz, D. García-Martín, S. Kazi, M. Larocca, and M. Cerezo, Showcasing a barren plateau theory beyond the dynamical lie algebra, arXiv preprint arXiv:2310.11505 (2023).

[20] S. Raj, I. Kerenidis, A. Shekhar, B. Wood, J. Dee, S. Chakrabarti, R. Chen, D. Herman, S. Hu, P. Minssen, *et al.*, Quantum deep hedging, Quantum **7**, 1191 (2023).

[21] L. Monbroussou, J. Landman, A. B. Grilo, R. Kukla, and E. Kashefi, Trainability and expressivity of hamming-weight preserving quantum circuits for machine learning, arXiv preprint arXiv:2309.15547 (2023).

[22] M. Levin and X.-G. Wen, Detecting topological order in a ground state wave function, Physical review letters **96**, 110405 (2006).

[23] A. Y. Kitaev, Quantum computations: algorithms and error correction, Russian Mathematical Surveys **52**, 1191 (1997).