Are Quantum Convolutional Neural Networks Actually Useful? Likely not!

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*

Framework: Quantum Convolutional Neural Networks (QCNN) [1] have seen widespread use in Quantum Machine Learning (QML) as they posses desirable features such as being implementable in logarithmic depth and not exhibiting barren plateaus [2]. In particular, these models have shown success for classifying quantum phases of matter [3–7], and classical data embedded in quantum states [5, 8, 9]. While some works have attempted to understand what these models can learn for a *given* problem [10], they do not address the elephant in the room: *Should we even use a QCNN at all*? In this work we present theoretical and numerical evidence which allows us to argue that:

- Randomly initialized QCNNs prevent barren plateaus by only being able to "see" their input states locally. Mathematically, this means that the Heisenberg evolved local measurement operator only has support on low-bodied Paulis. As such, they can only capture local information from the input states, and therefore their action during the first steps of training is effectively restricted to a classically simulable polynomial subspace.
- While one could argue that as the training progresses the QCNN action might drift to a non-classically simulable region, we find that for most tasks where QCNNs have shown heuristic success, they do so by staying in the polynomial subspace. That is, during the beginning, middle, and final steps of training the QCNN only sees the input state locally. As an example, we present in this abstract the task of classifying trivial states from topologically protected states in a bond-alternating Heisenberg model and a Haldane chain, benchmarks that have been widely used to showcase QCNNs. Yet, one can readily see that a trivial state and a topologically protected state look different locally (almost by definition), and therefore there exist local measurements which tell the states apart.
- The previous point then raises the concerning issue that most QCNN studies in the literature consider classification problems which can be fully solved with access to local information (i.e., they are fairly trivial and can also be tackled classically given local measurements from the input states). As such, we propose an adversarial supervised classification dataset where we test if a QCNN can classify between two topologically protected states. Importantly, all the states in the dataset now look maximally mixed when seen locally, so the QCNN must leave the classically simulable region to solve the task. Unfortunately, we find that in this case the QCNN over-fits the data and is unable to classify.
- Taken together, our results indicate that: (i) QCNNs appear to work only in tasks for which there exists classical algorithms that can also solve the problem when given access to data coming from a quantum computer, and (ii) when presented with a non-trivial task, QCNNs will be unable to solve it.

Importantly, we highlight that our work should not be understood as *definitively* proving that QCNNs are not useful. We see our results as only presenting sobering evidence regarding their performance in challenging datasets, and shedding serious doubts on their general usefulness. Moreover our work further strengthens the connection between absence of barren plateaus and classical simulability.

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 alternat-

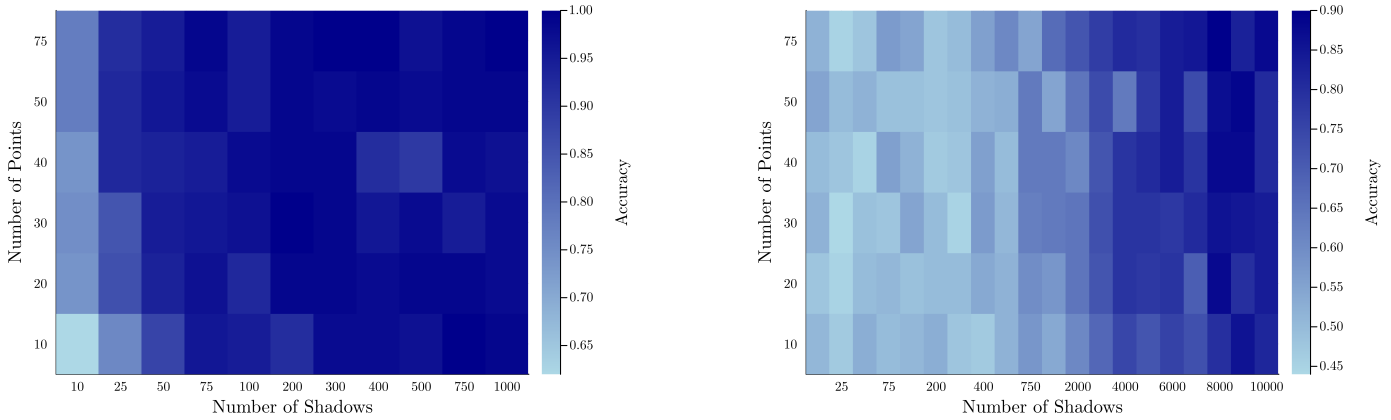

Figure 1: Classification accuracy for the simulation of a QCNN using LOWESA restricted to polynomial subspace operators and performing classical shadows on the input states. We note that the results shown here correspond to Bond-alternating model (left) and Haldane chain (right) with $n = 100$ and $n = 30$ qubits, respectively, for a set of 100 samples.

ing 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. 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 quantum phases of matter classification tasks [11]. However, by revisiting several results in the literature, we argue, and numerically support, that the following result 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.*

The key implication of Result 2 is that the QCNN should be classically simulable *once* we have a classical representation of the quantum data states in the subspace of local operators. Importantly, some form of tomographic procedure (e.g. collecting classical shadows [12–14]) is needed to get this data. This experiment could be performed on a quantum computer but could also potentially be performed on a simpler (e.g. non-universal, analogue) quantum simulator. These measurements are used to approximate the action of the QCNN on its initial states. More formally, we argue that for tasks where QCNNs work, then the algorithms is simulable in $\mathrm{CSIM}_{\mathrm{QE}}$ [15]. (A problem $\mathcal{C}$ is in $\mathrm{CSIM}_{\mathrm{QE}}$, i.e., is 'quantum-enhanced classical simulation', 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 claims, we classically simulate the action of a QCNN via tensor networks and the LOWESA algorithm [16] for the task of classifying between the 2 phases of the bond-alternating XXX Heisenberg model [17] and the 2 phases of the Haldane chain [18]. We recall that the ground

states of H_{XXX} and H_H are characterized by a trivial phase and a topologically protected phase. We consider a model of 100 qubits for H_{XXX} and 30 qubits for H_H, and employ a dataset of 100 samples, with equal distribution of both ground state species. To illustrate that the action of the QCNN needs not leave the space of local observables to accurately solve the task, we perform classical Pauli shadows on the data states, and we bias the action of the QCNN so that it can only explore the polynomially-sized operators subspace.

In Fig 1 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. First, one can see that the accuracy increases as the number of shadows increases. Then, one also observes the slight dependence of the accuracy on the number of training points, 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. 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. That this means that one can classically solve the task given the shadow tomography with much less calls to the quantum computer.

Next, we consider a binary classification task where the states is obtained by perturbing the degenerate ground states of the toric code and double semion realized on a torous; two topological states of matter with non-local nature of the order parameter between phases. Here we simulate a randomly initialized QCNN acting on the full state as input (i.e., no tomographic procedure, as we want to study the QCNN as it would be implemented on a quantum device). As we show in Fig 2, the QCNN over-fits the input states independently of the number of points and is not able to perform any training at all, evidencing the proof that a randomly initialized QCNN only "sees" a polynomial size operator subspace which does not include the *global* order parameter that achieves the correct classification. Therefore, for non-trivial tasks, which could potentially stand beyond classical numerical tools, QCNN seems incapable to function, restraining its potential use to a set of tasks which mostly fall into the classically simulable region.

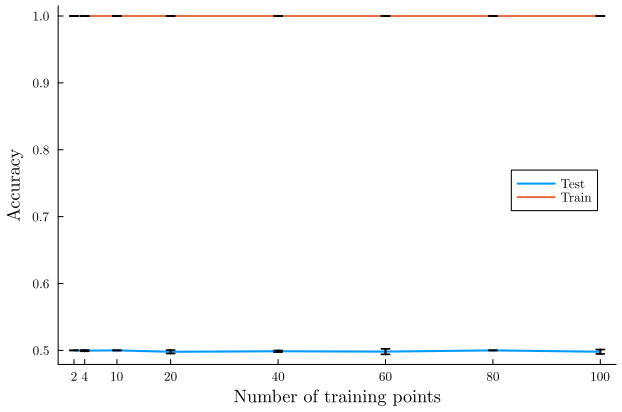

Figure 2: We show the accuracy of a QCNN for the classification of a dataset generated with topological ground states of the toric code and double semion model on a torus with n = 12 qubits. The total number of samples is 500, and the results are averaged over 5 independent runs.

Implications and Future Directions: At a fundamental level, our work shows that QCNNs solve tasks by only exploring the classically simulable region of local observables on the initial states. Hence, for the tasks where they have shown heuristic success [19–27], there appears to be no need for running a classifying model on a quantum computer as a classical algorithm with access to measurements can perform equally well. Indeed, our results reveal that many heuristic benchmarks performed in the literature to showcase the power of QML models could be misleading, and should be therefore revisited [28]. Our results leave open the possibility that there could be future realizations of QCNNs where gradient flows lead to regions of the parameter space where signals from global observables become important. However, we have found no evidence of that in our explorations so far. Thus, from a practical point of view our realizations are extremely positive, as training a classical model based on measurements (such as classical shadows) without running on-chip any variational algorithm with expensive optimization loops results in major savings in terms of quantum resources.

[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, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, Theory for equivariant quantum neural networks, PRX Quantum **5**, 020328 (2024).

[8] G. Chen, Q. Chen, S. Long, W. Zhu, Z. Yuan, and Y. Wu, Quantum convolutional neural network for image classification, Pattern Analysis and Applications **26**, 655 (2023).

[9] Y. Li, R.-G. Zhou, R. Xu, J. Luo, and W. Hu, A quantum deep convolutional neural network for image recognition, Quantum Science and Technology **5**, 044003 (2020).

[10] C. Umeano, A. E. Paine, V. E. Elfving, and O. Kyriienko, What can we learn from quantum convolutional neural networks?, arXiv preprint arXiv:2308.16664 (2023).

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

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

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

[14] 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).

[15] 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).

[16] E. Fontana, M. S. Rudolph, R. Duncan, I. Rungger, and C. Cîrstoiu, Classical simulations of noisy variational quantum circuits, arXiv preprint arXiv:2306.05400 (2023).

[17] A. Kitazawa, K. Nomura, and K. Okamoto, Phase diagram of $S = 1$ bond-alternating xxz chains, Phys. Rev. Lett. **76**, 4038 (1996).

[18] F. D. M. Haldane, Nonlinear field theory of large-spin Heisenberg antiferromagnets: semiclassically quantized solitons of the one-dimensional easy-axis néel state, Physical Review Letters **50**, 1153 (1983).

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

[20] 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).

[21] 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).

[22] 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).

[23] L. Schatzki, M. Larocca, Q. T. Nguyen, F. Sauvage, and M. Cerezo, Theoretical guarantees for permutation-equivariant quantum neural networks, npj Quantum Information **10**, 12 (2024).

[24] M. T. West, J. Heredge, M. Sevior, and M. Usman, Provably trainable rotationally equivariant quantum machine

learning, arXiv preprint arXiv:2311.05873 (2023).
[25] 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).
[26] 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).
[27] 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).
[28] J. Bowles, S. Ahmed, and M. Schuld, Better than classical? the subtle art of benchmarking quantum machine learning models, arXiv preprint arXiv:2403.07059 (2024).