

Hybrid quantum-classical neural network for sample-efficient recognition of topological phases

Markus K. Hoffmann,¹ Leon C. Sander,¹ Colin Scarato,^{2,3} Johannes Knörzer,^{2,3,4} Luca Hofele,^{2,3} Shengpu Wang,^{2,3} Kilian Hanke,^{2,3} Ashay Sathe,¹ Dominic Haggmann,^{2,3,4} Alexander Flasby,^{2,3,4} Andreas Wallraff,^{2,3,4} Christoph Hellings,^{2,3} Michael J. Hartmann,^{1,5,6} and Petr Zapletal¹

¹*Department of Physics, Friedrich-Alexander Universität Erlangen-Nürnberg (FAU), Erlangen, Germany*

²*Department of Physics, ETH Zurich, CH-8093 Zurich, Switzerland*

³*Quantum Center, ETH Zurich, CH-8093 Zurich, Switzerland*

⁴*ETH Zurich - PSI Quantum Computing Hub, Paul Scherrer Institute, CH-5232 Villigen, Switzerland*

⁵*Max Planck Institute for the Science of Light, Staudtstraße 2, 91058 Erlangen, Germany*

⁶*Quint Computing GmbH, Erlangen, Germany*

(Dated: July 8, 2026)

I. INTRODUCTION

Quantum computers can prepare many-body quantum states that cannot be fully described classically, making their characterization a key challenge. While randomized measurements such as classical shadows efficiently access many local observables, their sample complexity becomes prohibitively large for nonlocal correlations [1]. Alternatively, quantum states can be further processed on a quantum computer before measurement to enable more efficient characterization [2]. Here, we introduce a hybrid quantum-classical neural network, depicted in Fig. 1, combining a shallow parameterized quantum circuit, measurements and a classical feedforward neural network. The quantum circuit learns a nonlocal transformation of the measurement basis that enhances the distinguishability of different quantum states, while the classical network performs the final classification.

Classifying phases of matter is an important task in condensed matter physics and a promising near-term application for quantum computers [3]. Topological phases are particularly challenging to characterize because they lack local order parameters and are instead defined by inherently nonlocal features. Here, we demonstrate that the hybrid neural network distinguishes the topological phase of the surface code, depicted in Fig. 1a, not only from a topologically trivial phase, but also from a symmetry-enriched topological phase [4] and random product states forming a featureless 1-design. These datasets are deliberately designed to exclude local features. Successful classification therefore requires learning genuinely nonlocal features. The hybrid neural network reduces both training and inference sample complexities by approximately one order of magnitude compared to a purely classical neural network trained on randomized Pauli measurements.




FIG. 1. Hybrid neural network for recognizing the topological phase of the surface code. (a) Surface code realized by the tensor-network state for $g = 1$ and $N = 24$ qubits depicted as black points and crosses with the Hamiltonian $H_{SC} = -\sum_s A_s - \sum_p B_p$, where A_s and B_p are products of four Pauli X_i and Z_i operators shown as gray and green squares, respectively. (b) The ground states of the surface code are processed in a parameterized quantum circuit and (c) a classical feedforward neural network assigns an output $y(x)$ to each measured bit string $x = x_1 x_2 \dots x_N$.

Importantly, the hybrid neural network employs a much shallower quantum circuit than fully quantum neural networks [5]. This enabled us to experimentally realize it on a superconducting quantum processor and demonstrate that it is robust against decoherence, gate infidelities and state-preparation-and-measurement errors [6]. With clear working principles and sampling advantage shown in our theory work [7], and experimental realization reported in our experimental work [6], we establish the hybrid neural network as a promising framework for efficient characterization of complex quantum states on current quantum processors.

II. HYBRID QUANTUM-CLASSICAL NEURAL NETWORK

To recognize the phase of a quantum state $|\psi\rangle$, we process it in a hybrid neural network depicted in Fig. 1 consisting of a parameterized quantum circuit $U(\phi)$, measurements and a classical feedforward neural network, assigning an output $y(x) \in [0, 1]$ to each measured bit string x . The quantum circuit consists of parameterized single-qubit Pauli-Y rotations, and parameterized two-qubit controlled phase gates, which can be natively implemented on a superconducting quantum processor [6]. The classical feedforward neural network comprises a single fully-connected hidden layer with weights and biases w . We train the hybrid neural network using a two-loop procedure to minimize the binary cross-entropy cost with respect to the weights and biases w in the inner loop, and the quantum-circuit parameters ϕ in the outer loop.

The hybrid neural network has clear working principles. From the measured data, the classical neural network estimates the probability that a bit string x originates from the phase $L = 1$ as $y(x) \approx y_{\text{opt}}(x) = P_1(x)/[P_0(x) + P_1(x)]$ [8], where $P_L(x)$ denotes the measurement probability in phase L . Consequently, the optimal cost $C_{\text{opt}} = \ln 2 - D_{JS}[P_0, P_1]$ is determined by the Jensen-Shannon divergence $D_{JS}[P_0, P_1]$, a statistical distance between distributions P_0 and P_1 . Because characteristic nonlocal correlations of topological phases are hard to infer from direct measurements, we train the quantum circuit $U(\phi)$ to maximize the statistical distance $D_{JS}[P_0, P_1]$ by minimizing the cost with respect to the parameters ϕ . Thus, the hybrid neural network learns a measurement basis that enhances the distinguishability of the phases.

III. RECOGNITION OF TOPOLOGICAL PHASES AND ITS SAMPLE COMPLEXITY

We train the hybrid neural network on a family of tensor-network states introduced in Ref. [4] $|\psi\rangle = \sum_{x \in \{0,1\}^N} \text{tTr}[T^{x_1, x_2}(g), \dots, T^{x_{N-1}, x_N}(g)] |x\rangle$, where $|x\rangle = |x_1 x_2 \dots x_N\rangle$, $T(g)$ are rank-6 tensors parameterized by $g \in [-1, 1]$ and tTr denotes the tensor contraction. For $g > 0$, the states belong to the surface-code phase. For $g < 0$, they realize another topological phase, also featuring topological order, but enriched by




FIG. 2. Distinguishing the surface-code phase, $g > 0$, from the symmetry-enriched topological phase, $g < 0$. (a) Mean output $\bar{y}$ of the hybrid neural network for training states (circles) and test states (diamonds). (b) Multi-shot error rate $\bar{p}_{\text{err}}$ vs. the total number of training measurement shots S_{tot} for the hybrid neural network (blue diamonds) and a classical neural network trained on randomized measurements (orange diamonds).




FIG. 3. Experimental training for the surface-code phase and random product states. (a) Traces of the cost in the outer optimization loop for $N = 16$ and 10 random initializations. (b) Output of the trained network for $N = 9$ (pink) and $N = 16$ (blue) vs. magnetic field strength h . Vertical dashed lines indicate the phase boundary with a trivial magnetic phase.

an anti-unitary $\mathbb{Z}_2^T$ symmetry. Since both phases lack local order parameters, their classification requires the detection of topological features. Fig. 2a shows the output of the trained hybrid neural network as a function of g , successfully distinguishing the two phases also for states unseen during training (diamonds). Furthermore, we investigate in Fig. 2b the dependence of the classification error on the number of measurement shots S_{tot} used in training. The hybrid neural network achieves an error rate of $\bar{p}_{\text{err}} < 0.01\%$ using $S_{\text{tot}} \approx 10^6$ training shots.

We benchmark the hybrid neural network against a classical feedforward neural network with five fully-connected hidden layers and trained on randomized Pauli measurements. Fig. 2b shows that the classical network requires $S_{\text{tot}} \approx 10^7$ training shots to achieve the same error rate of $\bar{p}_{\text{err}} < 0.01\%$, approximately one order of magnitude larger than the number of measurements required for the hybrid network.

IV. EXPERIMENTAL DEMONSTRATION ON SUPERCONDUCTING HARDWARE

We implement the hybrid neural network on a superconducting quantum processor [6] and train it with quantum hardware in the loop using a second dataset comprising ground states of the surface code in a magnetic field and random product states forming a 1-design. Fig. 3a shows the cost as a function of training iteration of the quantum circuit. Fig. 3b shows the output of the trained network for the surface code perturbed by a magnetic field, successfully detecting the transition into a trivial magnetic phase. Furthermore, the trained network distinguishes the topological ground states of the surface code from random product states, even when perturbed by random single-qubit Pauli errors. We observe a classification accuracy above 85% for single-shot inference and above 99% when averaging over 10 measurement shots. Since the 1-design formed by random product states exhibits featureless measurement statistics, this successful classification shows that the network learned topological features of the surface code.

V. CONCLUSION

We demonstrate that the hybrid neural network reduces training and inference sample complexities by approximately one order of magnitude compared to the classical benchmark network trained on randomized measurements. Furthermore, we introduce two datasets that consist exclusively of phases lacking local order parameters and thus require the learning of genuinely topological features. In doing so, we address a challenge in quantum machine learning, where many existing benchmarks such as those in [5] have recently been classified with classical machine learning models using information encoded in local observables [9]. Crucially, we demonstrate its experimental viability by realizing the hybrid neural network on a superconducting quantum processor. Together, these results show that the hybrid quantum-classical neural network provides a practical route toward sample-efficient characterization of complex quantum states on near-term quantum hardware.

-
- [1] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
 - [2] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
 - [3] K. J. Satzinger, Y. Liu, A. Smith, C. Knapp, M. Newman, C. Jones, Z. Chen, C. Quintana, X. Mi, A. Dunsworth, *et al.*, Realizing topologically ordered states on a quantum processor, *Science* **374**, 1237 (2021).
 - [4] Y.-J. Liu, K. Shtengel, and F. Pollmann, Simulating two-dimensional topological quantum phase transitions on a digital quantum computer, *Phys. Rev. Res.* **6**, 043256 (2024).
 - [5] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).
 - [6] C. Scarato, J. Knörzer, M. K. Hoffmann, L. C. Sander, L. Hofele, S. Wang, K. Hanke, A. Sathe, D. Hagmann, A. Flasby, M. J. Hartmann, P. Zapletal, A. Wallraff, and C. Hellings, Hybrid quantum-classical neural networks for recognizing quantum phases, *arXiv preprint* (2026), [arXiv:2606.28201](#).
 - [7] M. K. Hoffmann, L. C. Sander, C. Scarato, C. Hellings, J. Knörzer, M. J. Hartmann, and P. Zapletal, Hybrid quantum-classical neural network for sample-efficient recognition of topological phases, *arXiv preprint* (2026), [arXiv:2606.28199](#).
 - [8] J. Arnold and F. Schäfer, Replacing neural networks by optimal analytical predictors for the detection of phase transitions, *Phys. Rev. X* **12**, 031044 (2022).
 - [9] P. Bermejo, P. Braccia, M. S. Rudolph, Z. Holmes, L. Cincio, and M. Cerezo, Quantum convolutional neural networks are effectively classically simulable, *PRX Quantum* **7**, 020304 (2026).