

Quantum data learning of topological-to-ferromagnetic phase transitions in the 2 + 1D toric code loop gas model

Shamminuj Aktar^{1,*}, Rishabh Bhardwaj¹, Andreas Bärtzsch¹, Tanmoy Bhattacharya², and Stephan Eidenbenz¹

¹*Computing and Artificial Intelligence Division (CAI-3), Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

²*Theoretical Division (T-2), Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*



(Received 11 December 2025; accepted 8 May 2026; published 26 June 2026)

Quantum data learning (QDL) provides a framework for extracting physical insights directly from quantum states, bypassing the need for any identification of the classical observable of the theory. A central challenge in many-body physics is that the identity of quantum phases, especially those with topological order, is often inaccessible through local observables or simple symmetry-breaking diagnostics. Here, we apply QDL techniques to the 2 + 1-dimensional toric-code loop-gas model in a magnetic field. Ground states are generated across multiple lattice sizes using a parametrized loop-gas circuit with a variational quantum-eigensolver approach. We then train a quantum convolutional neural network across the full field-parameter range to perform phase classification and capture the overall phase structure. We also employ a physics-aware training protocol that excludes the near-critical region ($0.2 \leq x \leq 0.4$) around $x_c \approx 0.25$, the phase-transition point estimated by quantum Monte Carlo simulation, reserving this window for testing to evaluate the ability of the model to learn the phase transition. In parallel, we implement an unsupervised quantum k -means method based on state overlaps, which partitions the dataset into two phases without prior labeling. Our supervised QDL approach recovers the phase structure and accurately locates the phase transition, in close agreement with previously reported values; the unsupervised QDL approach recovers the phase structure and locates the phase transition with a small offset as expected in finite volumes; both QDL methods outperform classical alternatives. These findings establish QDL as an effective framework for characterizing topological quantum matter, studying finite volume effects, and probing phase diagrams of higher-dimensional systems.

DOI: [10.1103/4d35-kw5d](https://doi.org/10.1103/4d35-kw5d)

I. INTRODUCTION

The study of quantum many-body systems lies at the heart of condensed-matter and high-energy physics [1]. Strong correlations among many particles lead to emergent behavior that cannot be reduced to single-particle descriptions, underpinning phenomena such as high-temperature superconductivity, quantum spin liquids, and topological order. Classical computational approaches, including density functional theory (DFT) [2,3], quantum Monte Carlo (QMC) [4,5], and the density matrix renormalization group [6,7], have provided profound insights into correlated quantum systems. However, each of these methods faces intrinsic limitations: DFT relies on approximate exchange-correlation functionals that break down in strongly

correlated regimes; QMC is constrained by the fermion sign problem and unstable analytic continuation; and tensor-network approaches are efficient only for low-entanglement states, with poor scaling in higher dimensions or real-time dynamics. More recently, classical machine learning has been leveraged for phase classification and order-parameter discovery in many-body systems [8–14]. Yet the exponential growth of Hilbert space, extensive postprocessing required, and absence of local order parameters remain major obstacles to studying large many-body systems. These limitations motivate the exploration of alternative strategies that exploit quantum computational resources. As argued by Huang *et al.* [15], quantum models can, in principle, efficiently extract physical information inaccessible to classical learners, suggesting a new path for studying correlated matter.

Quantum machine learning (QML), a rapidly developing framework, combines the expressive power of parametrized quantum circuits with the adaptive strategies of modern learning algorithms [16–20]. Most QML applications to date have focused on encoding classical data into quantum states, which are subsequently processed by hybrid

*Contact author: saktar@lanl.gov

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

quantum-classical models for classification or regression tasks [21–23]. This paradigm has enabled progress in diverse domains, including quantum chemistry [24], materials science [25], and high-energy physics [26–28]. However, it inherits the cost of designing efficient encodings and does not fully utilize the inherent quantum structure of many-body states. Recently, a new direction, *quantum data learning* (QDL), has been proposed, in which quantum models learn directly from quantum-mechanical data. This eliminates the need for classical encodings and enables direct manipulation of quantum wave functions.

QDL techniques are naturally suited for studying high-energy physics. Recent developments in collider experiments are enabling the direct probing of quantum states through measurable observables such as angular distributions, momentum correlations, and precision timing, offering access to entanglement entropy and quantum coherence [29–32]. Moreover, quantum tomography has already been demonstrated in systems of top quarks, photon pairs, and entangled mesons [33–36], together with quantum-information analyses in quantum chromodynamics [32]. With upcoming detector upgrades at the Large Hadron Collider expected to provide real-time access to such quantum observables, QDL offers a promising framework for extracting physical insights directly from high-dimensional quantum data.

Earlier QDL studies have focused on abstract or low-dimensional systems. Nagano *et al.* [37] employed quantum convolutional neural network (QCNN) architectures on $(1 + 1)$ -dimensional systems, including the Schwinger model and the $\mathbb{Z}_2$ lattice gauge theory, identifying quantum phases and extracting coupling parameters in parton-shower-inspired simulations. These studies demonstrated that QDL can uncover nontrivial structures without explicit order parameters, motivating extensions to higher-dimensional systems.

In this work, we investigate QDL techniques in the $2 + 1$ -dimensional toric-code model in a magnetic field, a paradigmatic system that hosts a transition between a topological phase and a ferromagnetic phase [38–42]. The toric code model captures the essential physics of confinement while remaining accessible to both analytical control and efficient simulation. It exhibits topological order, emergent anyonic excitations, and a confinement-deconfinement transition, making it a minimal yet rich platform for probing the interplay between quantum correlations and phase structure. From a practical standpoint, the model is simple enough to be tractable on classical computers and implementable on near-term quantum devices, while also serving as a meaningful stepping stone toward more complex gauge theories, including non-Abelian models and ultimately quantum chromodynamics. Its relevance extends beyond theory: recent advances in programmable quantum simulators, such as Rydberg-atom arrays, have realized quantum dimer models with nonlocal string operators, directly observing hallmarks of toric-code

physics including anyonic excitations, coherent topological sectors, and loop-parity measurements [43,44].

To probe the phase structure of the toric-code model, we generate ground states using a parametrized loop-gas circuit (PLGC) within a variational quantum eigensolver (VQE), following the approach of Sun *et al.* [42], for lattice sizes up to 4×3 . First, we train QCNNs across the full range of the field parameter and observe nearly perfect classification accuracy, demonstrating that phase information can be extracted directly from the quantum ground states. We then adopt a physics-aware strategy in which states far from the QMC-estimated transition at $x_c \approx 0.25$ [38,41] are used for training, while those in the region $(0.2 \leq x \leq 0.4)$ are reserved for testing. From the QCNN decision function, we extract phase transition as flip interval estimators x'_c for each lattice size and observe systematic convergence toward x_c as system size increases. In parallel, we implement an unsupervised quantum k -means algorithm based on fidelity overlaps, which partitions the data into two phases without labels. The resulting cluster flip interval estimators $\hat{x}_c$ capture the overall phase boundary and broadly follow the QCNN-predicted transition points x'_c , though with a small offset.

Finally, we benchmark QDL in learning phase transition of the toric-code model against classical machine-learning approaches, including convolutional neural networks (CNNs) and logistic regression. Both the CNN and logistic regression models were trained separately on the VQE ansatz parameters and on the quantum state amplitudes as input features. While classical models capture the broad phase structure, QCNNs consistently yield more reliable transition points, particularly for the 4×3 lattice where the estimate is approaching the QMC estimate. Moreover, when varying the training-set size, QCNNs are more sample-efficient than other baselines in learning the phase transition. Finite-size scaling further shows that the QCNN-extrapolated transition $x_c(\infty) = 0.234 \pm 0.022$ converges closely to the QMC benchmark $x_c \approx 0.25$.

The rest of this work is organized as follows. Section II outlines the ground-state generation using the parametrized loop-gas ansatz and describes the quantum data learning framework, including supervised QCNN and unsupervised quantum k -means models. Section III presents the experimental results on learning phase transition, comparing QDL-based phase learning with classical machine-learning methods and performing finite-size scaling. Finally, Sec. IV summarizes the findings and discusses future directions.

II. METHODOLOGY

Figure 1 provides an overview of the QDL framework applied to the $2 + 1$ -dimensional toric-code loop-gas model under a magnetic field. First, ground states are prepared within a VQE setup across different lattice geometries. The resulting states serve as quantum inputs $|\psi_x\rangle$ for the learning stage. These input states are then processed either




FIG. 1. Overview of the workflow for quantum data learning of phases in the 2 + 1-dimensional toric-code loop-gas model under a magnetic field. Ground states of $H(x)$ are approximated by a VQE using the PLGC on finite lattices. The resulting states $|\psi_{\text{in}}\rangle$ are processed using QDL techniques to learn the phase structure and distinguish between topologically ordered and ferromagnetically ordered phases.

by QCNNs for supervised phase-transition learning or by quantum k -means models for unsupervised clustering. Both approaches enable the identification of the transition between the topologically ordered and ferromagnetically ordered phases.

A. Toric code loop-gas model

Topological phases represent a remarkable class of quantum matter whose low-energy properties remain robust against local perturbations and cannot be characterized by conventional symmetry-breaking order parameters. These phases have attracted substantial interest in condensed-matter and quantum-information science due to their inherent stability, which makes them promising platforms for fault-tolerant quantum computation [38]. A paradigmatic example of a topologically ordered system is Kitaev's toric code, an exactly solvable lattice model described by the Hamiltonian [45]

$$H_{\text{TCM}} = -\sum_s A_s - \sum_p B_p, \quad (1)$$

where $A_s = \prod_{i \in s} \sigma_i^x$ (star operators acting on edges adjacent to vertex s) and $B_p = \prod_{i \in p} \sigma_i^z$ (plaquette operators acting on edges surrounding plaquette p) mutually commute. Figure 1 illustrates these operators on a square lattice, highlighting a star operator A_s (red) and a plaquette operator B_p (blue). The ground state of the toric code is a coherent superposition of closed-loop spin configurations defined on the lattice edges, commonly referred to as a quantum loop gas.

Under periodic boundary conditions (PBCs), opposite edges of the lattice are identified, mapping the system onto a torus T^2 . This topology eliminates boundary effects and associated edge excitations (anyons), leaving only intrinsic topological properties. A key consequence is a degenerate

ground-state manifold composed of distinct topological sectors. These sectors correspond to elements of the first homology group $H_1(T^2, \mathbb{Z}_2)$ and are characterized by the $\mathbb{Z}_2$ winding parities along noncontractible loops in the x and y directions. Specifically, the winding sectors are defined by Wilson-loop operators

$$W_{x/y}^a = \prod_{i \in c_{x/y}} \sigma_i^a,$$

where $c_{x/y}$ denotes noncontractible cuts along the x/y directions, and the index $a = x, z$ labels the electric and magnetic sectors, respectively. This construction is illustrated schematically in Fig. 2 (left). The ground states obtained by coherent superpositions of these sectors are immune to mixing by any local perturbations.

When open boundary conditions (OBCs) are applied, the lattice effectively reduces to a disk shown in Fig. 2 (right). As a result, the topology-protected ground-state degeneracy is lost. The bulk topological order, however, persists: loop configurations behave as in the toric case, spins remain long-range entangled, and the system supports deconfined anyons whose boundaries can absorb or condense them. The phase retains a nonzero topological entanglement entropy $\gamma = \ln 2$ [46]. Depending on the boundary termination ("rough" or "smooth"), edge anyon condensation lifts the global degeneracy, while in geometries with multiple boundaries (e.g., a cylinder or punctured plane) a reduced, boundary-dependent degeneracy reemerges. In essence, OBCs on finite systems remove the global degeneracy and introduce anyonic edge modes but do not destroy the underlying $\mathbb{Z}_2$ topological order. In the present work, we employ OBCs, which allow for explicit treatment of edge effects and simplify the implementation on finite lattices by removing nonlocal stabilizer constraints and periodic wraparound couplings required under PBCs.




FIG. 2. Toric geometry and boundary effects in the toric code model (TCM). Left: construction of the toric geometry under periodic boundary conditions, yielding a manifold with two noncontractible Wilson loops W_x and W_y that define distinct topological sectors of the ground state and encode the intrinsic $\mathbb{Z}_2$ topological order. Right: under open boundary conditions, the finite square lattice is topologically equivalent to a disk D^2 , eliminating noncontractible loops and hence the global ground-state degeneracy while preserving the bulk $\mathbb{Z}_2$ order.

External perturbations such as a uniform magnetic field introduce an effective loop tension that destabilizes the topological phase. When the field is sufficiently strong, it selects a single, topologically trivial sector and drives a continuous second-order quantum phase transition into a ferromagnetically ordered state [38–42]. Physically, this transition corresponds to the condensation of magnetic vortices and the confinement of emergent $\mathbb{Z}_2$ charges, providing a canonical example of the breakdown of topological order.

Following the methodology introduced by Sun *et al.* [42], we study the toric-code model in a uniform longitudinal magnetic field. Qubits are placed on the lattice edges under OBC, and the Hamiltonian interpolates continuously between the topologically ordered and ferromagnetic limits,

$$H(x) = (1 - x)H_{\text{TCM}} - x \sum_{i=1}^N \sigma_i^z, \quad (2)$$

where the parameter $x \in [0, 1]$ tunes the relative strength between the stabilizer terms (A_s, B_p), which favor topological order at $x = 0$, and the magnetic field term (σ_i^z), which favors ferromagnetic alignment at $x = 1$. Quantum Monte Carlo simulations estimate the critical transition point between these phases at $x_c \approx 0.25$ in the thermodynamic limit [38,41]. Traditional QMC analyses identify the critical point using a local order parameter, typically the average magnetization along the z axis,

$$\langle m_z \rangle = \frac{1}{N} \left\langle \sum_{i=1}^N \sigma_i^z \right\rangle, \quad (3)$$

and by examining the Binder cumulant

$$Q(m_z) = \frac{\langle m_z^2 \rangle^2}{\langle m_z^4 \rangle},$$

together with its finite-size scaling behavior near x_c . In contrast, the present work aims to predict this critical point

by learning directly from the ground states prepared on a quantum processor.

B. Toric-code ground-state preparation

To study quantum phases and their transitions, it is essential to obtain accurate approximations of the ground states of the system Hamiltonian. In principle, such ground states can be prepared using quantum phase estimation [47] or computed via exact diagonalization (ED); however, both approaches are resource intensive and become impractical on hardware available today even for modest lattice sizes. In this work, we instead employ the VQE, which offers a hardware-efficient framework for approximating ground states.

We adopt the physics-motivated PLGC introduced in Ref. [42]. The construction begins from the trivial product state $|0\rangle^{\otimes N}$ and applies one unitary per plaquette, replacing the Hadamard gate in the exact toric-code preparation with a tunable rotation $R_y(\theta_p)$. The resulting variational state takes the form

$$|\psi_x(\theta)\rangle = \prod_{p=1}^{N_p} \left[\cos\left(\frac{\theta_p}{2}\right) I + \sin\left(\frac{\theta_p}{2}\right) B_p \right] |0\rangle^{\otimes N}, \quad (4)$$

where $B_p = \prod_{i \in p} \sigma_i^x$ is the plaquette operator acting on the four edges surrounding plaquette p . This circuit generates weighted superpositions of closed-loop configurations, interpolating between the equal-weight loop gas at $\theta_p = \pi/2$ and the trivial product states at $\theta_p = 0$ or π . The PLGC thus confines the variational search space within the physically relevant loop-gas subspace while retaining tunability across different field regimes [42].

For each lattice size, optimized ground states are prepared across a grid of field values x . The VQE cost function is the energy expectation value

$$E(\theta; x) = \langle \psi(\theta) | H_{\text{TCM}}(x) | \psi(\theta) \rangle, \quad (5)$$

which is estimated from quantum measurements on the PLGC ansatz. Details of the VQE training procedure are provided in Appendix A. After generating the ground

states, we assign phase labels to the states based on our approximate knowledge of the phase transition point: states with $x \leq x_c$ are identified with the topologically ordered phase (-1) , while those with $x > x_c$ correspond to the ferromagnetically ordered phase $(+1)$, where $x_c \approx 0.25$ is the QMC-estimated critical point in the thermodynamic limit [38,41]. In Sec. II C 1, we describe a quantum data learning method to refine this input approximate x_c to a precise finite-volume transition point, and in Sec. II C 2, we devise an alternate method to find the transition point without using any such input.

C. Quantum data learning framework

1. Supervised learning: QCNN

For supervised phase classification, we employ QCNNs, a hierarchical circuit ansatz inspired by multiscale entanglement renormalization [48]. QCNNs possess a logarithmic-depth architecture that is provably resilient to barren plateaus, enabling efficient training on near-term quantum devices [49]. In addition, the number of variational parameters grows only polylogarithmically with system size, and the architecture is naturally compatible with fault-tolerant implementations. Previous studies have demonstrated that QCNNs are powerful tools for quantum phase recognition: Cong *et al.* first proposed QCNNs for identifying topological phases in spin chains [48] and subsequent works have extended this framework for other systems [50–52]. Recently, Nagano *et al.* applied them to lattice gauge theories and high-energy physics simulations [37], successfully identifying confinement, deconfinement, and coupling parameters directly from quantum states. These successes motivate their application to the toric-code model considered here.

Following the construction in Ref. [37], our QCNN alternates convolution and pooling layers in a recursive manner, reducing the number of qubits at each stage while retaining global information about the system (see Appendix B for the QCNN architecture diagram and the circuit implementations of each convolution and pooling block). The ground states $|\psi_x\rangle$ prepared by the VQE serve directly as inputs to the QCNN. The first convolution layer consists of two-qubit $SU(4)$ unitaries arranged in a brickwork pattern across the full register, efficiently entangling local degrees of freedom. The subsequent pooling layers employ parametrized single-qubit rotations followed by CNOT gates to project out qubits, effectively renormalizing the system while preserving some nonlocal correlations. The following convolution layers act on fewer qubits by pairing only adjacent ones, and this coarse-graining process continues until only two qubits remain before the final pooling step. The last pooling block outputs a single qubit, which is measured in the Z basis to yield the QCNN output.

The resulting measurement outcome serves as a phase classifier, distinguishing the $\mathbb{Z}_2$ topologically ordered phase from the ferromagnetically ordered phase. Formally, the

QCNN prediction for an input quantum state $|\psi_x\rangle$ is defined as the expectation value of the Pauli- Z operator on the final output qubit after processing through the QCNN unitary circuit $U_{\text{QCNN}}(\phi)$,

$$y_{\text{out}}(\phi) = \text{Tr}[Z_{\text{out}} U_{\text{QCNN}}(\phi) |\psi_x\rangle \langle \psi_x| U_{\text{QCNN}}^\dagger(\phi)], \quad (6)$$

where Z_{out} acts on the remaining qubit, and ϕ denotes the set of trainable parameters in the QCNN. The output $y_{\text{out}} \in [-1, 1]$ serves as a continuous phase indicator: $y_{\text{out}} \leq 0$ corresponds to the topologically ordered phase (-1) , while $y_{\text{out}} > 0$ corresponds to the ferromagnetically ordered phase $(+1)$.

To identify the phase boundary, we analyze the QCNN decision function as a function of the tuning parameter x . We employ a robust estimator of the transition point by defining x'_c as the central value of the interval $[x_{j_{\min}}, x_{j_{\max}}]$, where $j_{\min}$ is the largest index such that all samples with $x_i \leq x_{j_{\min}}$ share the same QCNN predicted phase label, and $j_{\max}$ is the smallest index such that all samples with $x_i \geq x_{j_{\max}}$ share the opposite phase label. The estimated flip interval point is then

$$x'_c = \frac{1}{2}(x_{j_{\min}} + x_{j_{\max}}), \quad (7)$$

with an uncertainty of $\Delta x'_c = (x_{j_{\max}} - x_{j_{\min}})/2$ associated to it.

2. Unsupervised learning: Quantum k means

In parallel, we explore an unsupervised approach based on quantum k -means clustering that relies on state overlaps [53,54]. Given a set of ground states $\{|\psi_i\rangle\}$, pairwise similarities are quantified by the state fidelity,

$$F(\psi_i, \psi_j) = |\langle \psi_i | \psi_j \rangle|^2. \quad (8)$$

From this, we define the Hilbert-Schmidt distance [55] between pure states as

$$d(\psi_i, \psi_j) = \sqrt{2(1 - F(\psi_i, \psi_j))}, \quad (9)$$

which provides a distance metric on the Hilbert space. Clustering is performed with $k = 2$ by minimizing the total intracluster distance,

$$\mathcal{L} = \sum_{c=1}^k \sum_{i \in C_c} d^2(\psi_i, \psi_{m_c}), \quad (10)$$

where ψ_{m_c} denotes the medoid state of cluster C_c . Cluster labels are oriented such that small- x states correspond to the topological phase, ensuring consistent interpretation across system sizes. To identify the phase boundary, we analyze the cluster assignments as a function of the tuning

parameter x , similar to the QCNN approach. The estimated cluster flip interval point is then

$$\hat{x}_c = \frac{1}{2}(x_{j_{\min}} + x_{j_{\max}}), \quad (11)$$

with an associated uncertainty calculated identically to that used for the QCNN. Unlike the QCNN, quantum k -means requires no labeled data (and hence requiring no prior knowledge of the transition point even approximately) and partitions the quantum states directly into phases based solely on their pairwise overlaps, providing an unsupervised estimate of the transition point.

III. EXPERIMENTAL RESULTS

A. Toric-code dataset: Ground states from VQE

Using the procedure described in Sec. II B, we generate datasets of size 300 for lattices of size 2×2 , 2×3 , 3×3 , and 4×3 , corresponding to $N = 4, 7, 12$, and 17 qubits, respectively. Each dataset consists of 300 quantum ground states, with sampling divided evenly between the two phases: 150 states are prepared at equidistant field values $x \in [0, x_c]$, covering the topologically ordered regime, and the remaining 150 at $x \in [x_c + 10^{-2}, 1]$, covering the

ferromagnetically ordered regime. This yields balanced datasets with uniform coverage of the parameter range.

To validate the accuracy of the generated states, we compare the VQE-optimized ground states against ED results [56]. The comparison includes the ground-state energy per qubit as

$$E_N = \frac{1}{N} \langle \psi(\theta) | H_{\text{TCM}}(x) | \psi(\theta) \rangle, \quad (12)$$

and the average magnetization per qubit as

$$\langle m_z \rangle_N = \frac{1}{N} \sum_{i=1}^N \langle \sigma_i^z \rangle. \quad (13)$$

Appendix A shows that the per-qubit ground-state energies and average magnetizations obtained from VQE and exact diagonalization agree almost perfectly. Following Sec. II B, each ground state is assigned a phase label of (-1) for the topologically ordered phase and $(+1)$ for the ferromagnetically ordered phase.

B. Phase learning via QCNN

To learn phase classification, we first train the QCNN model described in Sec. II C 1 using a random 80%–20%




FIG. 3. QCNN evaluation with randomly split datasets. Each panel corresponds to a different lattice size: 2×2 , 2×3 , 3×3 , and 4×3 . Predicted phase labels closely follow the true labels for all lattice sizes, demonstrating near-perfect classification accuracy across system sizes.

train-test split across the full range of the field parameter x . For each lattice size, 80% of the toric-code states are used for training while the remaining 20% are reserved for testing. The QCNN is optimized using the Adam optimizer with an initial learning rate of 0.01. Training is performed for 100 epochs with a batch size of 24 and includes an L_2 regularization penalty with strength $\lambda = 10^{-4}$ to suppress overfitting. More details on the training procedure and the learning curves are shown in Appendix C. After training, we evaluate the model on the testing states to compute the classification accuracy.

As shown in Fig. 3, the predicted phase labels closely follow the true labels across all lattice sizes. The model achieves near-perfect classification accuracy: 100% for the 2×2 , 2×3 , and 4×3 lattices, and 96% for the 3×3 lattice. Each panel of Fig. 3 corresponds to a specific lattice size, allowing direct comparison across system sizes and demonstrating excellent agreement between QCNN-predicted and labeled phases.

To investigate the phase-transition region, we adopt a physics-aware data split based on the QMC-estimated infinite-volume critical point $x_c \approx 0.25$ [38,41]. In this setting, ground states away from the transition ($x < 0.2$ or $x > 0.4$) are used for training, while those within the

boundary window ($0.2 \leq x \leq 0.4$) are reserved for testing. This choice of boundary window ensures approximately balanced representations of both phases in the training and testing datasets. The training setup is same as random-split training discussed earlier. For each lattice size, a QCNN is trained on the off-critical data and evaluated on the test region. From the predicted phases in the boundary region, we also compute phase flip interval estimators x'_c as discussed in Sec. II C 1.

The raw QCNN outputs as a function of x in the region ($0.2 \leq x \leq 0.4$) are shown in Fig. 4, where each subplot corresponds to one of the four lattice sizes (2×2 , 2×3 , 3×3 , and 4×3). The vertical yellow dashed lines indicate the phase flip interval points x'_c , along with their range of variance, extracted from the QCNN decision boundary, while the red dashed line marks the infinite-volume QMC critical point $x_c \approx 0.25$ for reference.

For the 2×2 lattice, the estimate is $x'_c = 0.272(2)$, which decreases to $x'_c = 0.267(2)$ for 2×3 . The 3×3 system yields $x'_c = 0.282(2)$, while the 4×3 lattice produces $x'_c = 0.246(1)$, in close agreement with the infinite-volume QMC value. Further study is needed to understand the nonmonotonic finite-volume effects implied by these observations.




FIG. 4. QCNN phase-boundary evaluation across lattice sizes. Raw QCNN outputs $\langle Z \rangle$ are shown as functions of the field parameter x for lattices 2×2 , 2×3 , 3×3 , and 4×3 . The yellow-shaded regions indicate the phase flip-interval estimates of the transition point $x'_c \pm \Delta x$ extracted from the QCNN decision boundary, while the red dashed lines mark the QMC benchmark $x_c \approx 0.25$. As the system size increases, the QCNN-estimated phase boundary becomes close to the infinite-volume QMC critical point.




FIG. 5. Quantum k -means phase classification of toric-code ground states. Each panel corresponds to a different lattice size: 2×2 , 2×3 , 3×3 , and 4×3 . The shaded yellow regions denote the phase transition intervals $\hat{x}_c' \pm \Delta x$, while the red dashed line marks the QMC reference $x_c \approx 0.25$. The cluster assignments (+1 and -1) (jittered for visibility) exhibit clear phase separation and finite-size effects similar to the QCNN results shown in Fig. 4.

C. Phase detection via quantum k -means

We next apply the unsupervised quantum k -means clustering procedure described in Sec. II C 2 to the toric-code loop-gas datasets for each lattice size. Using the distance measure based on state overlaps as described in Eq. (10), the algorithm partitions the data into two clusters that are observed to correspond to the topological and ferromagnetic phases, without requiring any prior label information. We compute $\hat{x}_c$ as cluster flip interval point as discussed in Sec. II C 2. Figure 5 shows the resulting phase clusters for all lattice sizes, together with the QMC reference point x_c and the cluster-flip estimates $\hat{x}_c$.

The estimated cluster-flip points $\hat{x}_c$ are 0.262(2), 0.272(2), 0.282(2), and 0.277(2) for the 2×2 , 2×3 , 3×3 , and 4×3 lattices, respectively, are similar to the QCNN phase flip estimates x_c' . This demonstrates that the phase structure can be extracted directly from quantum data in an unsupervised manner, perhaps with slight overestimation. Although in this small study the QCNN aligns more closely with the QMC benchmark, especially for the 4×3 lattice, the quantum k -means approach remains a strong alternative given that no labels or prior knowledge are required.

D. Comparison with classical machine-learning models

To benchmark the QDL approaches against classical counterparts, we train CNNs and logistic regression classifiers on the same toric-code loop-gas datasets. Both classical models take as input either the state amplitudes $|\psi_x|^2$ in the computational basis or the corresponding VQE ansatz parameters θ . The dimensionality of the amplitude vectors grows exponentially with the number of qubits, becoming computationally intractable for larger lattice sizes. Using the variational parameters θ as input is also not ideal, since these parameters are not directly accessible for an arbitrary ground state. Previous studies have explored classical shadows or sampled representations as alternative inputs, but such approaches require additional postprocessing [15,57–60]. Since the QCNN operates directly on quantum states, our comparison focuses exclusively on models trained with either state amplitudes or VQE ansatz parameters as input.

We compare the phase flip-interval estimations obtained from the QCNN, quantum k -means, CNN, and logistic regression models. Each model is trained on randomly selected subsets of 50, 100, 200, and 300 samples drawn from the off-critical regions of the magnetic-field parameter




FIG. 6. Top: (left) QCNN x'_c ; (right) quantum k -means $\hat{x}_c$. Middle: CNN (left) and logistic regression (right) trained on quantum-state amplitudes $|\psi|^2$. Bottom: CNN (left) and logistic regression (right) trained on PLGC parameters θ (parameter-based models). Each subplot shows the estimated phase interval point as a function of training-sample size for lattices 2×2 , 2×3 , 3×3 , and 4×3 . For better visibility, a small horizontal offset is applied to the data points along the x axis. Classical models exhibit systematic shifts from the QMC benchmark ($x_c \approx 0.25$) of a size comparable to the shift seen in our QCNN experiment. The QCNN converges toward x_c for larger lattices even with 200 training samples while the quantum k -means model closely tracks the phase boundary without supervision.

x , while reserving 300 samples within the boundary window ($0.2 \leq x \leq 0.4$) for testing. This setup ensures that all models are trained on balanced phases within sampled states and are evaluated on the total unseen states near the phase transition. Each training procedure is repeated ten times, and the resulting phase flip-interval estimates are reported as the mean and standard deviation across repetitions.

Figure 6 presents a comparison of the average phase flip-interval estimations between the quantum and classical models. The top panels show the QCNN estimates x'_c (left) and the quantum k -means estimates $\hat{x}_c$ (right). The middle panels display the estimated x'_c for the CNN and logistic regression models trained on quantum-state amplitudes $|\psi|^2$, while the bottom panels correspond to CNN and logistic regression models trained on the PLGC parameters θ .




FIG. 7. Finite-size scaling of estimated critical points $x'_c(L)$ versus $1/L_{\text{eff}}$ for QCNN, quantum k -means, CNN, and logistic regression. Solid lines denote weighted fits to the scaling form $x'_c(L) = x_c(\infty) + a/L_{\text{eff}}$, and markers indicate results for each lattice size. Error bars on the data points represent the flip-interval uncertainty $\Delta x'_c$ (for QDL methods) or the standard deviation across repeated training runs (for classical models); error bars on the extrapolated points at $1/L_{\text{eff}} \rightarrow 0$ represent the one-sigma fitting uncertainty. The red dashed line shows the QMC reference $x_c \approx 0.25$.

Each subplot presents the estimated transition point as a function of training-sample size for lattice geometries up to 4×3 .

For the QCNN, the phase estimates converge toward $x_c \approx 0.25$ for larger lattice sizes with relatively few samples compared to the classical models. The 3×3 system exhibits a slight upward shift in x'_c , while the 4×3 lattice converges near the QMC benchmark with as few as 100 samples. For the quantum k -means algorithm, the cluster-flip estimates $\hat{x}_c$ remain slightly above the QMC value (0.26–0.29) yet stable with increasing sample size. In contrast, both CNN and logistic regression models show less consistent and more erratic learning behavior for both amplitude-based and VQE-parameter-based inputs. Their phase flip estimates also systematically overestimate the transition point relative to the QMC reference. We further investigate both classical models and find that their predictions fluctuate within the testing region $0.2 \leq x \leq 0.4$. Appendix C shows the phase-label prediction plots for both classical models considered in this comparison in Fig. 12. We also find that classical models trained on PLGC inputs exhibit smaller fluctuations in phase prediction than those trained on amplitudes. This indicates that the PLGC representation encodes more robust phase-related information than the amplitude representation. Overall, we conclude that the QDL models demonstrate the ability to learn phase boundaries directly from quantum states without requiring order parameters or prior labeling as well as or better than the classical approaches.

E. Finite-size scaling

To consolidate the results across different lattice sizes, we perform a finite-size scaling analysis of the estimated phase-transition points $x'_c(L)$. These estimates are obtained from QDL techniques, including the QCNN and quantum k means, as well as from classical models (CNN and logistic regression) trained on either VQE ansatz parameters or amplitude representations. Statistical uncertainties on each data point are propagated from the flip-interval width $\Delta x'_c$ (for QDL methods) or the standard deviation across repeated training runs (for classical models), and are shown as error bars in Fig. 7.

For each method, the critical-point estimates are fitted to the empirical scaling relation

$$x'_c(L) = x_c(\infty) + \frac{a}{L_{\text{eff}}}, \quad (14)$$

where $L_{\text{eff}} = \sqrt{p}$ with p the number of plaquettes, and $x_c(\infty)$ denotes the extrapolated critical point in the thermodynamic limit. The $1/L_{\text{eff}}$ correction follows the standard leading-order finite-size scaling form near a continuous transition [38,41]. $L_{\text{eff}} = \sqrt{p}$ provides a uniform effective linear scale across the nonsquare lattice geometries used in this study. A standard χ^2 -minimizing fit with each data point weighted by the inverse-squared error produces an unacceptably large χ^2 per degree-of-freedom ($\chi^2/\text{d.o.f.}$), possibly because the errors on each individual point are being underestimated. In this process, to obtain an

estimate of the error on the extrapolated value, we follow the standard procedure of inflating the error on the individual points by a constant factor so as to obtain the modified $\chi^2/\text{d.o.f.} = 1$. The resulting fits are shown in Fig. 7. Though Akaike Information Criterion favors inclusion of an $O((a/L_{\text{eff}})^2)$ correction, we do not do so for a few reasons: (i) the $\chi^2/\text{d.o.f.}$ fails to improve on inclusion of the additional term, (ii) the extrapolated values agree within errors, but the extrapolation errors increase significantly, (iii) the scatter of points suggests underestimated errors rather than a uniform curvature as the source of the large χ^2 , and (iv) for such small lattices, the asymmetry might control the finite size scaling corrections more than the quadratic terms, and a full analysis of that effect is beyond the scope of this paper.

Even though the 3×3 lattice exhibits an anomalously deviant finite-size effect, the QCNN yields an extrapolated critical point of $x_c(\infty) = 0.234 \pm 0.022$, which lies within one standard deviation of the QMC benchmark $x_c \approx 0.25$. We note that with only four modest lattice sizes available, the finite-size scaling fits are not highly constrained, and the extrapolated values should be regarded as indicative estimates of the thermodynamic limit rather than precise determinations.

The quantum k -means method gives $x_c(\infty) = 0.2930 \pm 0.0060$, slightly overestimating the transition, consistent with the small upward shift observed in the finite-volume cluster assignments. Among the classical approaches, the CNN trained on PLGC inputs extrapolates to $x_c(\infty) = 0.318 \pm 0.013$, while the amplitude-based CNN gives $x_c(\infty) = 0.278 \pm 0.002$. Logistic regression yields $x_c(\infty) = 0.2903 \pm 0.0013$ for PLGC inputs and $x_c(\infty) = 0.338 \pm 0.028$ for amplitude inputs. Both CNN models and the amplitude-based logistic regression exhibit significant deviations from the QMC benchmark, with LogReg-Amplitude showing the largest overestimation among all methods. The PLGC-based logistic regression performs comparably to the quantum k means in terms of the extrapolated value, but shows less reliable convergence behavior, as reflected in the scatter of data points visible in Fig. 7.

Overall, the QCNN provides the most accurate and stable estimation, closely reproducing the thermodynamic transition point within its one-sigma confidence range and showing convergence toward the QMC benchmark as system size increases. More studies are necessary to fully characterize the quality of the various finite-size extrapolation schemes.

IV. CONCLUSION

Experimental results demonstrate that QDL provides an effective framework for probing the topological-to-ferromagnetic phase transition in the $2 + 1$ -dimensional toric-code loop-gas model. QCNNs achieve near-perfect phase classification and yield phase flip estimates that converge toward the QMC benchmark, with finite-size scaling giving $x_c(\infty) = 0.234 \pm 0.022$ in excellent

agreement with the previously known result $x_c \approx 0.25$. Remarkably, QCNNs exhibit robust scaling behavior even when trained on data from only four modest lattice sizes, underscoring their ability to extract critical behavior from limited quantum resources. The quantum k -means clustering method similarly identifies the phase boundary without requiring labeled data, albeit with a small upward shift, while classical CNN and logistic-regression models consistently overestimate the transition point.

The main limitation of the present study lies in the modest lattice sizes (up to 4×3) accessible by classical simulation, where finite-size fluctuations remain non-negligible. Future extensions will aim to scale QDL approaches to larger lattices such as [61,62] and more complex models with intrinsic topological order such as [63,64], supported by improved state-preparation techniques and hardware-efficient implementations on near-term quantum devices.

Building on these results, our QDL study of the toric-code transition suggests a promising route toward learning phase diagrams of lattice gauge theories relevant to quantum chromodynamics (QCD). The QCD phase diagram features confinement-deconfinement and chiral transitions, potential critical endpoints, and intricate finite-density and finite-temperature behavior that remain difficult to access with classical methods [65]. As an intermediate, controllable step, a $\mathbb{Z}_2$ lattice gauge theory with dynamical fermions [64] provides a natural testbed for QDL-based classifiers. Its phase structure—encompassing confined “hadronic” regimes, deconfined or quark-liquid-like sectors, and superfluid or Mott-insulating phases depending on the coupling and chemical potential—shares qualitative features with QCD and is well suited for exploration using supervised QCNNs and unsupervised, overlap-based clustering methods aimed at detecting first-order transitions, crossovers, and critical points.

Extending beyond the ground-state transitions explored here, the study of thermal transitions in such gauge-matter systems will require the preparation of mixed quantum states. Recently developed thermal pure quantum methods for gauge theories [63] enable efficient construction of finite- T and finite- μ states on quantum hardware, providing a natural foundation for applying QDL to finite-temperature and finite-density physics. Together, these directions outline a broader program for employing QDL architectures to map the phase structure of strongly correlated gauge theories and to bridge quantum simulation, quantum information, and high-energy physics.

ACKNOWLEDGMENTS

The research presented in this article was supported by the NNSA’s Advanced Simulation and Computing Beyond Moore’s Law program at Los Alamos National Laboratory and the Laboratory Directed Research and Development program of Los Alamos National Laboratory under Project

No. 20260043DR. This work has been assigned LANL technical report number LA-UR-25-30864.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available because of legal restrictions preventing unrestricted public distribution. The data are available from the authors upon reasonable request.

APPENDIX A: VQE TRAINING

We employed the PLGC ansatz within the VQE framework to generate toric-code ground states for $x \in [0, 1]$. Optimization is performed using the simultaneous perturbation stochastic approximation algorithm [66] with a learning rate of 0.01. Each VQE run was initialized with random angles $\theta_p \in [0, 2\pi)$, and multiple independent trials (ten per x value) were performed to mitigate convergence to local minima. Each trial proceeded for 500 optimization steps, and the parameters yielding the lowest final energy





FIG. 8. Schematic QCNN network topology for the 2×3 lattice. Alternating convolution (CL) and pooling (PL) layers progressively reduce the number of active qubits, followed by a final measurement on the last qubit.



were retained. We computed the ground-state energy per qubit and average magnetization per qubit, $\langle m_z \rangle$, defined in Eq. (3), for all lattice sizes.

Figure 9 compares the ground states obtained from ED with those generated by VQE across different lattice sizes. The left panel shows that the energy per qubit from VQE (symbols) is nearly indistinguishable from the ED results (solid lines) over the full field range, while the right panel shows similarly close agreement in the average magnetization. Both observables capture the transition from the low-field topological phase to the high-field ferromagnetic phase. These results suggest that the PLGC ansatz accurately reproduces the ground-state physics of the toric-code model under a magnetic field, even for larger lattices where ED becomes computationally expensive.

APPENDIX B: QCNN MODEL ARCHITECTURE

The QCNN model used in this work consists of alternating CL and PL layers designed to hierarchically coarse-grain quantum information encoded in the toric-code ground states [37]. Figure 8 illustrates the QCNN architecture, which comprises seven qubits shown in Fig. 1 corresponding to the 2×3 lattice. Each convolution layer applies a set of parametrized two-qubit $SU(4)$ gates arranged in a brickwork pattern, followed by pooling operations that progressively reduce the number of active qubits from seven to four, then to two, and finally to one. The CL blocks perform local entangling transformations that extract multiqubit correlations and emulate the action of short-range renormalization operators, while the PL blocks implement nonunitary coarse graining through measurement-inspired controlled rotations that transfer relevant information to the control qubits while discarding redundant degrees of freedom. Figure 10 shows the internal




FIG. 9. Comparison between exact diagonalization (ED, solid lines) and VQE simulations using the PLGC for lattice sizes 2×2 , 2×3 , 3×3 , and 4×3 . (a) Ground-state energy per qubit E/N as a function of the tuning parameter x . (b) Average magnetization per qubit $\langle m_z \rangle = N^{-1} \sum_i \langle \sigma_i^z \rangle$. The close agreement demonstrates the accuracy of VQE-generated states in reproducing both energy and order parameters across system sizes.




FIG. 10. QCNN building blocks. (a) CL implementing an $SU(4)$ -like two-qubit unit. (b) PL transferring correlations to the control qubit and discarding redundant degrees of freedom.

circuits of the convolution block (top) and the pooling block (bottom), parametrized by the learnable parameters $\{\alpha_i\}$ and $\{\beta_i\}$.

APPENDIX C: LEARNING PHASE TRANSITION

We train both QCNN and CNN models to learn phase classification and to estimate the phase-transition region. Each model is trained on a randomly selected 80% subset of ground states from the dataset and evaluated on the remaining 20%. Figure 11 (top) shows the training losses for both the QCNN (left) and the CNN models trained on either VQE ansatz parameters (middle) or state amplitudes

(right). The plots indicate that the CNN models converge faster and achieve lower training losses than the QCNN, reflecting the simpler optimization landscape of classical architectures.

The QCNN is optimized using the Adam algorithm with an initial learning rate of 0.01, batch size of 24, and an L_2 regularization strength $\lambda = 10^{-4}$ to suppress overfitting. The cost function used for training is the binary cross-entropy between predicted probabilities and target phase labels, augmented by an L_2 penalty on the QCNN parameters,

$$\mathcal{L} = -\frac{1}{N} \sum_i [y_i \log p_i + (1 - y_i) \log(1 - p_i)] + \lambda \frac{\|\phi\|_2^2}{N}, \quad (\text{C1})$$

where $p_i = (1 + y_{\text{out}}(\phi))/2$ is the predicted probability derived from the QCNN output $y_{\text{out}}(\phi) = \langle Z_{\text{out}} \rangle_i$ as defined in Eq. (6), and $y_i \in \{0, 1\}$ denotes the target phase label. The first term measures the misclassification error, while the L_2 term regularizes the trainable parameters ϕ to improve generalization and prevent overfitting.

Training is performed for 100 epochs, during which the learning rate is adaptively decayed according to a step schedule to ensure stable convergence. Optimization is




FIG. 11. Comparison of training losses for QCNN and CNN models. Top row: training losses using random-split data for QCNN (left), CNN with PLGC input (middle), and CNN with amplitude input (right). Bottom row: training losses using the physics-aware split excluding the near-critical region ($0.2 \leq x \leq 0.4$) for QCNN (left), CNN with PLGC input (middle), and CNN with amplitude input (right). The physics-aware split yields consistently lower losses and smoother convergence, as data outside the critical window are more easily separable into distinct phases. The CNN converges more rapidly toward zero loss, reflecting its simpler classical optimization landscape compared with the QCNN.




FIG. 12. Phase-label predictions (+1 and -1) from classical machine-learning models (jittered for visibility) trained on toric-code quantum-state datasets for the 4×3 lattice. The top panel shows phase predictions from CNNs using input PLGC parameters θ (left) and state amplitudes $|\psi\rangle$ (right). The bottom panel shows phase predictions from logistic-regression models using the same inputs, PLGC parameters θ (left) and amplitudes $|\psi\rangle$ (right). The red dashed line denotes the QMC reference ($x_c \approx 0.25$), while the yellow dashed line marks the flip interval estimator x'_c extracted from each model. Models trained on PLGC inputs yield more accurate phase boundaries with a smaller range of variance than those trained on amplitudes, although both deviate noticeably from the QMC benchmark.

monitored through the mean loss per epoch, and convergence is reached when successive loss values vary by less than 10^{-3} . The training curves in Fig. 11 illustrate this behavior.

To estimate the phase transition, each model is trained on ground states away from the critical region ($x < 0.2$ or $x > 0.4$) and evaluated within the testing region ($0.2 \leq x \leq 0.4$). As shown in Fig. 11 (bottom), the physics-aware split yields lower losses and smoother convergence than the random split, since training data in the critical window are potentially corrupted by finite volume effects. For this split, the CNN models also demonstrate faster convergence than the QCNN.

For both classical and quantum models, the phase transition point is determined from the model-predicted phase labels. Figure 12 shows phase-label predictions from classical machine-learning models: CNNs (top) and logistic regression (bottom) for the 4×3 lattice. The results show that models trained on VQE ansatz parameters (left) yield more accurate phase estimations than those trained on state amplitudes (right). In contrast, both CNN and logistic-regression models trained on amplitudes exhibit fluctuating and less stable predictions within the testing region. Although the loss curves in Fig. 11 indicate better convergence for the CNN models on training states, they still show poor generalization around the phase boundary.

- [1] A. Di Meglio *et al.*, Quantum computing for high-energy physics: State of the art and challenges, *PRX Quantum* **5**, 037001 (2024).
- [2] P. Hohenberg and W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* **136**, B864 (1964).
- [3] W. Kohn, Nobel lecture: Electronic structure of matter—wave functions and density functionals, *Rev. Mod. Phys.* **71**, 1253 (1999).
- [4] D. Ceperley and B. Alder, Quantum Monte Carlo, *Science* **231**, 555 (1986).
- [5] F. Becca and S. Sorella, *Quantum Monte Carlo Approaches for Correlated Systems* (Cambridge University Press, Cambridge, England, 2017).
- [6] S. R. White, Density matrix formulation for quantum renormalization groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [7] S. R. White, Density-matrix algorithms for quantum renormalization groups, *Phys. Rev. B* **48**, 10345 (1993).
- [8] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [9] J. Carrasquilla and R. G. Melko, Machine learning for quantum matter, *Quantum Sci. Technol.* **5**, 014002 (2020).
- [10] D.-L. Deng, X. Li, and S. Das Sarma, Machine learning topological states, *Phys. Rev. B* **96**, 195145 (2017).
- [11] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [12] L. Wang, Discovering phase transitions with unsupervised learning, *Phys. Rev. B* **94**, 195105 (2016).
- [13] J. F. Rodriguez-Nieva and M. S. Scheurer, Identifying topological order through unsupervised machine learning, *Nat. Phys.* **15**, 790 (2019).
- [14] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, Neural-network quantum states, string-bond states, and chiral topological states, *Phys. Rev. X* **8**, 011006 (2018).
- [15] 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).
- [16] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [17] M. Schuld and N. Killoran, Quantum machine learning in feature Hilbert spaces, *Phys. Rev. Lett.* **122**, 040504 (2019).
- [18] S. Mangini, F. Tacchino, D. Gerace, D. Bajoni, and C. Macchiavello, Quantum computing models for artificial neural networks, *Europhys. Lett.* **134**, 10002 (2021).
- [19] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, Challenges and opportunities in quantum machine learning, *Nat. Comput. Sci.* **2**, 567 (2022).
- [20] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [21] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature (London)* **567**, 209 (2019).
- [22] M. Schuld, Supervised quantum machine learning models are kernel methods, *arXiv:2101.11020*.
- [23] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).
- [24] O. A. von Lilienfeld, K.-R. Müller, and A. Tkatchenko, Exploring chemical compound space with quantum-based machine learning, *Nat. Rev. Chem.* **4**, 347 (2020).
- [25] M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, Fast and accurate modeling of molecular atomization energies with machine learning, *Phys. Rev. Lett.* **108**, 058301 (2012).
- [26] A. Blance and M. Spannowsky, Quantum machine learning for particle physics using a variational quantum classifier, *J. High Energy Phys.* **02** (2021) 212.
- [27] W. Guan, G. Perdue, A. Pesah, M. Schuld, K. Terashi, S. Vallecorsa, and J.-R. Vlimant, Quantum machine learning in high energy physics, *Mach. Learn.* **2**, 011003 (2021).
- [28] S. L. Wu and S. Yoo, Challenges and opportunities in quantum machine learning for high-energy physics, *Nat. Rev. Phys.* **4**, 143 (2022).
- [29] M. Fabbrichesi, R. Floreanini, and G. Panizzo, Testing Bell inequalities at the LHC with top-quark pairs, *Phys. Rev. Lett.* **127**, 161801 (2021).
- [30] C. Severi, C. D. E. Boschi, F. Maltoni, and M. Sioli, Quantum tops at the LHC: From entanglement to Bell inequalities, *Eur. Phys. J. C* **82**, 285 (2022).
- [31] J. A. Aguilar-Saavedra and J. A. Casas, Improved tests of entanglement and Bell inequalities with LHC tops, *Eur. Phys. J. C* **82**, 666 (2022).
- [32] M. Hentschinski, D. E. Kharzeev, K. Kutak, and Z. Tu, QCD evolution of entanglement entropy, *Rep. Prog. Phys.* **87**, 120501 (2024).
- [33] Y. Afik and J. R. M. de Nova, Entanglement and quantum tomography with top quarks at the LHC, *Eur. Phys. J. Plus* **136**, 907 (2021).
- [34] R. Aoude, E. Madge, F. Maltoni, and L. Mantani, Quantum SMEFT tomography: Top quark pair production at the LHC, *Phys. Rev. D* **106**, 055007 (2022).
- [35] M. Fabbrichesi, R. Floreanini, and E. Gabrielli, Constraining new physics in entangled two-qubit systems: Top-quark, tau-lepton and photon pairs, *Eur. Phys. J. C* **83**, 162 (2023).
- [36] K. Ehatäht, M. Fabbrichesi, L. Marzola, and C. Veelken, Probing entanglement and testing Bell inequality violation with $e^+e^- \rightarrow \tau^+\tau^-$ at Belle II, *Phys. Rev. D* **109**, 032005 (2024).
- [37] L. Nagano, A. Miessen, T. Onodera, I. Tavernelli, F. Tacchino, and K. Terashi, Quantum data learning for quantum simulations in high-energy physics, *Phys. Rev. Res.* **5**, 043250 (2023).
- [38] S. Trebst, M. Troyer, Z. Wang, and A. W. W. Ludwig, Breakdown of a topological phase: Quantum phase transition in a loop gas model with tension, *Phys. Rev. Lett.* **98**, 070602 (2007).
- [39] J. Vidal, S. Dusuel, and K. P. Schmidt, Low-energy effective theory of the toric code model in a parallel magnetic field, *Phys. Rev. B* **79**, 033109 (2009).

- [40] S. Dusuel, M. Kamfor, R. Orús, K. P. Schmidt, and J. Vidal, Robustness of a perturbed topological phase, *Phys. Rev. Lett.* **106**, 107203 (2011).
- [41] F. Wu, Y. Deng, and N. Prokof'ev, Phase diagram of the toric code model in a parallel magnetic field, *Phys. Rev. B* **85**, 195104 (2012).
- [42] R.-Y. Sun, T. Shirakawa, and S. Yunoki, Parametrized quantum circuit for weight-adjustable quantum loop gas, *Phys. Rev. B* **107**, L041109 (2023).
- [43] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin, Probing topological spin liquids on a programmable quantum simulator, *Science* **374**, 1242 (2021).
- [44] H. Weimer, M. Müller, I. Lesanovsky, P. Zoller, and H. P. Büchler, A Rydberg quantum simulator, *Nat. Phys.* **6**, 382–388 (2010).
- [45] A. Kitaev and C. Laumann, Topological phases and quantum computation, in *Exact Methods in Low-Dimensional Statistical Physics and Quantum Computing* (Oxford University Press, Oxford, 2010), Vol. 101.
- [46] This quantity is a topological invariant and is independent of lattice geometry.
- [47] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2000).
- [48] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).
- [49] A. Pesah, M. Cerezo, S. Wang, T. Volkoff, A. T. Sornborger, and P. J. Coles, Absence of barren plateaus in quantum convolutional neural networks, *Phys. Rev. X* **11**, 041011 (2021).
- [50] J. Herrmann, S. M. Lima, 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, *Nat. Commun.* **13**, 4144 (2022).
- [51] Y.-J. Liu, A. Smith, M. Knap, and F. Pollmann, Model-independent learning of quantum phases of matter with quantum convolutional neural networks, *Phys. Rev. Lett.* **130**, 220603 (2023).
- [52] S. Monaco, O. Kiss, A. Mandarino, S. Vallecorsa, and M. Grossi, Quantum phase detection generalization from marginal quantum neural network models, *Phys. Rev. B* **107**, L081105 (2023).
- [53] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum algorithms for supervised and unsupervised machine learning, [arXiv:1307.0411](https://arxiv.org/abs/1307.0411).
- [54] I. Kerenidis, J. Landman, A. Luongo, and A. Prakash, q -means: A quantum algorithm for unsupervised machine learning, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2019), Vol. 32.
- [55] V. V. Dodonov, O. V. Man'ko, V. I. Man'ko, and A. Wünsche, Hilbert-Schmidt distance and non-classicality of states in quantum optics, *J. Mod. Opt.* **47**, 633 (2000).
- [56] P. Weinberg and M. Bukov, QuSpin: A Python package for dynamics and exact diagonalisation of quantum many body systems part I: Spin chains, *SciPost Phys.* **2**, 003 (2017).
- [57] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [58] H.-Y. Huang, Learning quantum states from their classical shadows, *Nat. Rev. Phys.* **4**, 81 (2022).
- [59] G. Struchalin, Y. A. Zagorovskii, E. Kovlakov, S. Straupe, and S. Kulik, Experimental estimation of quantum state properties from classical shadows, *PRX Quantum* **2**, 010307 (2021).
- [60] T. Zhang, J. Sun, X.-X. Fang, X.-M. Zhang, X. Yuan, and H. Lu, Experimental quantum state measurement with classical shadows, *Phys. Rev. Lett.* **127**, 200501 (2021).
- [61] R. X. Siew, S. Chandrasekharan, and T. Bhattacharya, Asymptotic-freedom and massive glueballs in a qubit-regularized SU(2) gauge theory, [arXiv:2512.11068](https://arxiv.org/abs/2512.11068).
- [62] R. X. Siew, S. Chandrasekharan, and T. Bhattacharya, Continuum limit of a qubit-regularized SU(3) lattice gauge theory with glueballs, [arXiv:2603.01215](https://arxiv.org/abs/2603.01215).
- [63] Z. Davoudi, N. Mueller, and C. Powers, Towards quantum computing phase diagrams of gauge theories with thermal pure quantum states, *Phys. Rev. Lett.* **131**, 081901 (2023).
- [64] L. Homeier, A. Bohrdt, S. Linsel, E. Demler, J. C. Halimeh, and F. Grusdt, Realistic scheme for quantum simulation of $\mathbb{Z}_2$ lattice gauge theories with dynamical matter in $(2 + 1)d$, *Commun. Phys.* **6**, 127 (2023).
- [65] G. Aarts *et al.*, Phase transitions in particle physics, *Prog. Part. Nucl. Phys.* **133**, 104070 (2023).
- [66] J. Spall, Multivariate stochastic approximation using a simultaneous perturbation gradient approximation, *IEEE Trans. Autom. Control* **37**, 332 (1992).