

Quantum Phase Classification of Rydberg Atom Systems Using Resource-Efficient Variational Quantum Circuits and Classical Shadows

Hemish Ahuja*, Samradh Bhardwaj†, Kirti Dhir‡, Roman Bagdasarian§, and Ziwoong Jang¶

*†York University, Toronto, Canada

‡Modern School Vaishali, India

§Johannes Kepler Universität Linz, Linz, Austria

¶Pusan National University, Busan, South Korea

Abstract—Quantum phase transitions in Rydberg atom arrays present significant opportunities for studying many-body physics, yet distinguishing between different ordered phases without explicit order parameters remains challenging. We present a resource-efficient quantum machine learning approach combining classical shadow tomography with variational quantum circuits (VQCs) for binary phase classification of Z_2 and Z_3 ordered phases. Our pipeline processes 500 randomized measurements per 51-atom chain state, reconstructs shadow operators, performs PCA dimensionality reduction ($51 \rightarrow 4$ features), and encodes features using angle embedding onto a 2-qubit parameterized circuit. The circuit employs RY-RZ angle encoding, strong entanglement via all-to-all CZ gates, and a minimal 2-parameter ansatz achieving depth 7. Training via simultaneous perturbation stochastic approximation (SPSA) with hinge loss converged in 120 iterations. The model successfully learned to separate the classes in the limited dataset, achieving 100% accuracy on the training, validation, and held-out test sets. This work establishes pathways for quantum-enhanced condensed matter physics on near-term quantum devices.

Index Terms—quantum machine learning, Rydberg atoms, phase classification, variational quantum circuits, classical shadows, SPSA optimization

I. INTRODUCTION

Quantum phases of matter represent distinct organizational patterns characterized by spontaneous symmetry breaking and long-range quantum correlations [1]. Understanding and identifying these phases is fundamental to condensed matter physics and quantum simulation. Recent experimental advances in programmable Rydberg atom arrays have enabled the study of exotic quantum phases with unprecedented control [2], [3], yet classifying these phases without explicit order parameters remains a significant challenge.

A. Background and Motivation

Rydberg atoms, with their strong van der Waals interactions ($V(r) \sim C_6/r^6$) and controllable Rydberg blockade effect, have emerged as a versatile platform for quantum simulation [4]. In one-dimensional chains, these systems exhibit rich phase diagrams including Z_2 (period-2 ordering) and Z_3 (period-3 ordering) phases depending on the detuning δ and blockade radius R_b parameters. Traditional phase classification methods rely on computing order parameters specific to each

phase, requiring detailed knowledge of the system’s ground state—information that is exponentially expensive to obtain via full quantum state tomography.

The challenge intensifies in realistic experimental scenarios where only limited measurements are available. Full tomography of an N -qubit system requires $O(4^N)$ measurements, which is infeasible for even moderate system sizes. For a 51-atom Rydberg chain, this corresponds to approximately 10^{30} measurements—far beyond experimental capabilities.

B. Problem Statement

Given measurement outcomes from a 51-qubit Rydberg atom chain in an unknown quantum phase, with only 500 measurements per state in randomized Pauli bases, our goal is to perform binary classification: determine whether the system is in Z_2 or Z_3 ordered phase using minimal quantum circuit resources suitable for near-term quantum hardware. The constraints include: (1) no direct access to ground state wavefunction, (2) limited measurement budget, and (3) NISQ device limitations requiring shallow, narrow circuits.

C. Contributions

This work makes the following novel contributions:

1) Efficient Preprocessing Pipeline: We develop a classical shadow reconstruction protocol combined with PCA dimensionality reduction and angle-based feature encoding, reducing the problem from 51 qubits to 4 encoded features while preserving phase signatures.

2) Minimal Circuit Architecture: We design a 2-qubit variational circuit with depth 7 and only 2 trainable parameters, incorporating strong entanglement for correlation capture while avoiding barren plateau problems.

3) Hardware-Compatible Optimization: We implement SPSA gradient-free optimization with hinge loss for maximum-margin classification, demonstrating rapid convergence and robustness.

4) Perfect Classification: We achieve 100% test accuracy (3/3 samples), 100% validation accuracy (3/3 samples), and 100% training accuracy (14/14 samples), with a resource efficiency score of 0.5986.

II. THEORETICAL BACKGROUND

A. Rydberg Atom Systems

The Rydberg Hamiltonian for a one-dimensional chain of N atoms is given by:

$$H = \frac{\Omega}{2} \sum_{i=1}^N X_i - \delta \sum_{i=1}^N n_i + \sum_{i < j} \frac{\Omega R_b^6}{(a|i-j|)^6} n_i n_j \quad (1)$$

where Ω is the Rabi frequency driving transitions between ground $|g\rangle$ and Rydberg $|r\rangle$ states, δ is the detuning parameter, $n_i = |r\rangle\langle r|_i$ is the Rydberg number operator, R_b is the blockade radius, and a is the lattice spacing.

The first term represents coherent Rabi coupling, the second term penalizes Rydberg excitations proportional to detuning, and the third term captures van der Waals interactions between Rydberg atoms. The Rydberg blockade effect prevents multiple atoms within distance R_b from being simultaneously excited, fundamentally shaping the system's phase diagram.

1) Phase Diagram and Ordering Patterns: The system exhibits distinct quantum phases characterized by spatial ordering of Rydberg excitations:

Z₂ Phase (Period-2): Ground state approximately $|\dots rgrgrg \dots\rangle$ with alternating Rydberg excitations, breaking translational symmetry by 2 sites.

Z₃ Phase (Period-3): Ground state approximately $|\dots rggrggrg \dots\rangle$ with Rydberg excitations separated by 2 ground-state atoms, breaking translational symmetry by 3 sites.

Phase transitions occur by varying δ and R_b . Our dataset samples 20 points across this parameter space (10 Z₂, 10 Z₃), each prepared in the ground state and measured.

B. Classical Shadow Tomography

Classical shadows, introduced by Huang et al. [5], enable efficient estimation of quantum state properties using logarithmically fewer measurements than full tomography.

1) Shadow Generation Protocol: For each measurement $t = 1, \dots, T$:

- 1) Choose random Pauli basis: U_t (measurement in X , Y , or Z eigenbasis)
- 2) Measure in computational basis: outcome b_t
- 3) Compute shadow operator:

$$\sigma_t = 3U_t^\dagger |b_t\rangle\langle b_t| U_t - I \quad (2)$$

The classical shadow estimate of density matrix ρ is:

$$\hat{\rho} = \frac{1}{T} \sum_{t=1}^T \frac{\sigma_t + I}{3} \quad (3)$$

This provides an unbiased estimator: $\mathbb{E}[\hat{\rho}] = \rho$.

2) Measurement Basis Mapping: Our data provides measurements in Pauli eigenstates: $\{g, r, +, -, +i, -i\}$ corresponding to:

$$|g\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad |r\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (4)$$

$$|+\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad |-\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix} \quad (5)$$

$$|+i\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}, \quad |-i\rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \end{pmatrix} \quad (6)$$

These are eigenstates of Pauli Z , X , and Y operators respectively. For each outcome ψ , we construct:

$$\sigma(\psi) = 3|\psi\rangle\langle\psi| - I_2 \quad (7)$$

3) Advantage Over Full Tomography: Classical shadows require $O(\log N)$ measurements to estimate properties of N -qubit states with bounded error, compared to $O(4^N)$ for full tomography. For our 51-qubit system, 500 measurements suffice to capture phase-distinguishing features—a $\sim 10^{27}$ -fold reduction in measurement complexity.

C. Variational Quantum Circuits

Variational quantum circuits are parametrized quantum gates $U(\theta)$ that prepare quantum states $|\psi(\theta)\rangle = U(\theta)|0\rangle$ for hybrid quantum-classical optimization. The circuit comprises:

Feature Encoding: Maps classical data x to quantum state **Ansatz:** Parametrized gates with trainable weights θ **Measurement:** Extracts classical information (e.g., $\langle Z \rangle$)

A cost function $C(\theta)$ is minimized via classical optimizer updating θ based on quantum circuit evaluations. Key challenges include barren plateau problems (vanishing gradients in deep circuits) [6] and the expressibility-trainability tradeoff [7].

For phase classification, VQCs can capture quantum correlations in a sample-efficient manner, potentially providing advantages over classical machine learning when quantum features are relevant.

III. METHODOLOGY

A. Data Processing Pipeline

Our preprocessing pipeline transforms raw measurement data into quantum-encodable features through seven steps:

1) Step 1-3: Shadow Operator Construction: We load 20 data points, each containing 500 measurements across 51 qubits. For each qubit and measurement, we map the outcome to its corresponding basis state and construct the shadow operator via Eq. (2). Averaging over 500 measurements yields the classical shadow S_i for each qubit i :

$$S_i = \frac{1}{500} \sum_{t=1}^{500} \sigma_t^{(i)} \quad (8)$$

2) Step 4: Density Matrix Reconstruction: We reconstruct single-qubit density matrices:

$$\hat{\rho}_i = \frac{S_i + I_2}{3} \quad (9)$$

3) *Step 5: Pauli Expectation Values:* We compute Pauli X and Z expectation values:

$$\langle \sigma_x \rangle_i = \text{Re}[\text{Tr}(\hat{\rho}_i \sigma_x)] \quad (10)$$

$$\langle \sigma_z \rangle_i = \text{Re}[\text{Tr}(\hat{\rho}_i \sigma_z)] \quad (11)$$

where $\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ and $\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$.

This produces $20 \times 51 \times 2 = 2040$ real-valued features characterizing Bloch sphere coordinates for each qubit.

4) *Step 6-7: Angle Transformation and PCA:* We transform expectation values to angles:

$$\theta_x^{(i)} = (\text{clip}(\langle \sigma_x \rangle_i, -1, 1) + 1) \cdot \frac{\pi}{2} \quad (12)$$

$$\theta_z^{(i)} = (\text{clip}(\langle \sigma_z \rangle_i, -1, 1) + 1) \cdot \frac{\pi}{2} \quad (13)$$

Forming complex features $X_c = \theta_x + i\theta_z$ and separating real/imaginary parts yields a 20×102 feature matrix. Applying PCA with 4 components reduces dimensionality while retaining $> 90\%$ variance:

$$X_{\text{PCA}} = \text{PCA}([\text{Re}(X_c), \text{Im}(X_c)], n = 4) \quad (14)$$

Finally, we normalize to $[0, \pi]$ for angle encoding:

$$X_{\text{norm}} = \frac{X_{\text{PCA}} - X_{\min}}{X_{\max} - X_{\min}} \cdot \pi \quad (15)$$

Justification: PCA extracts latent structure corresponding to phase signatures. Z_2 and Z_3 phases have distinct spatial correlations that manifest as patterns in Pauli expectation values. PCA identifies directions of maximum variance, which correspond to phase-distinguishing collective modes.

B. Quantum Circuit Architecture

1) *Overall Structure:* Our circuit applies three layers to the initial state $|00\rangle$:

$$U(\mathbf{f}, \mathbf{w}) = U_{\text{ansatz}}(\mathbf{w}) \cdot U_{\text{ent}} \cdot U_{\text{encode}}(\mathbf{f}) \quad (16)$$

where $\mathbf{f} = [f_0, f_1, f_2, f_3]$ are input features and $\mathbf{w} = [w_0, w_1]$ are trainable parameters.

2) *Encoding Layer:* On qubit 0: $RY(f_0) \cdot RZ(f_1)$
On qubit 1: $RY(f_2) \cdot RZ(f_3)$

where rotation gates are:

$$RY(\theta) = \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{pmatrix} \quad (17)$$

$$RZ(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix} \quad (18)$$

This encoding maps 4 continuous features to arbitrary single-qubit states on the Bloch sphere. RY-RZ parameterization is universal for single-qubit rotations and provides efficient amplitude and phase encoding.

3) *Strong Entanglement Layer:* We apply controlled-Z (CZ) gates between all qubit pairs:

$$\text{CZ}_{01} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \quad (19)$$

For 2 qubits, this is a single CZ gate. Strong entanglement creates quantum correlations necessary for capturing spatial phase patterns. Since Z_2 and Z_3 phases exhibit long-range order, entanglement enables the circuit to detect these correlations.

4) *Variational Ansatz:* **Layer 1:** $RX(w_0)$ on both qubits, where

$$RX(\theta) = \begin{pmatrix} \cos(\theta/2) & -i\sin(\theta/2) \\ -i\sin(\theta/2) & \cos(\theta/2) \end{pmatrix} \quad (20)$$

Layer 2: CZ gate between qubits (ring topology)

Layer 3: $RZ(w_1)$ on both qubits

The shared parameters across qubits enforce structure and prevent overfitting. The RX-CZ-RZ sequence provides sufficient expressivity for binary classification while maintaining trainability.

5) *Circuit Metrics:*

- **Total Depth:** 7 gates (4 encoding + 1 entanglement + 2 ansatz)
- **Width:** 2 qubits
- **Parameters:** 2 (minimal)
- **Connectivity:** All-to-all (1 CZ pair)

6) *Measurement and Classification:* We measure all qubits in the computational basis and compute average Pauli-Z expectation:

$$\langle Z \rangle = \frac{1}{2} \sum_{i=0}^1 \langle Z_i \rangle \quad (21)$$

Classification rule: $\hat{y} = \text{sign}(\langle Z \rangle)$ where $Z_2 \rightarrow -1$ and $Z_3 \rightarrow +1$.

C. Training Protocol

1) *SPSA Optimization:* Simultaneous Perturbation Stochastic Approximation (SPSA) [8] approximates gradients using only 2 function evaluations per iteration regardless of parameter dimension:

Algorithm 1 SPSA Training

- 1: Initialize: $\theta_0 \sim \mathcal{N}(0, 0.01^2)$
 - 2: **for** $k = 0$ to $K - 1$ **do**
 - 3: Sample: $\delta_k \sim \text{Uniform}(\{-1, +1\}^2)$
 - 4: Compute: $c_k = c_0 / (k + 1)^\gamma$
 - 5: Evaluate: $L_+ = L(\theta_k + c_k \delta_k)$
 - 6: Evaluate: $L_- = L(\theta_k - c_k \delta_k)$
 - 7: Gradient: $\nabla L \approx \frac{L_+ - L_-}{2c_k \delta_k}$
 - 8: Update: $\theta_{k+1} = \theta_k - \alpha \nabla L$
 - 9: **end for**
-

Hyperparameters:

- Learning rate: $\alpha = 0.50$
- Initial perturbation: $c_0 = 0.40$
- Decay rate: $\gamma = 0.02$
- Maximum iterations: $K = 120$
- Adaptive shots: 256 (epochs 0-49), 512 (epochs 50-119)

SPSA is gradient-free, making it suitable for noisy quantum hardware where exact gradient computation is expensive. The $O(1)$ function evaluations per step (vs. $O(d)$ for parameter-shift rule) provide significant efficiency gains.

2) *Hinge Loss Function*: For binary classification with labels $y \in \{-1, +1\}$:

$$L_{\text{hinge}}(\mathbf{w}, X, \mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \max(0, 1 - y_i \langle Z \rangle_i(\mathbf{w})) \quad (22)$$

Hinge loss encourages confident predictions ($|y \cdot \langle Z \rangle| > 1$), creating a margin between classes similar to Support Vector Machines. This promotes generalization by penalizing samples within the margin even if correctly classified.

3) *Data Split*: From 20 total samples:

- Training: 14 samples (stratified: 7 Z_2 , 7 Z_3)
- Validation: 3 samples (stratified split)
- Test: 3 samples (stratified: 2 Z_2 , 1 Z_3)

Stratification ensures class balance across splits, critical for small datasets.

IV. RESULTS

A. Training Performance

Training converged rapidly, achieving 100% training accuracy by epoch 2 (iteration 40). Table I shows metrics recorded every 20 iterations.

TABLE I
TRAINING METRICS PER EPOCH

Epoch	Train Loss	Val Loss	Train Acc	Precision	F1
1	1.5573	1.3548	0.00%	0.00%	0.00%
2	0.3652	0.5833	100.00%	100.00%	100.00%
3	0.3744	0.5801	92.86%	100.00%	92.31%
4	0.3676	0.5814	100.00%	100.00%	100.00%
5	0.3693	0.5801	92.86%	100.00%	92.31%
6	0.3608	0.5651	100.00%	100.00%	100.00%
7	0.3708	0.6016	100.00%	100.00%	100.00%

The rapid convergence indicates a linearly separable feature space after PCA transformation. Small fluctuations in epochs 3 and 5 stem from SPSA's stochastic nature. Validation accuracy consistently matched training accuracy, showing no overfitting despite the small dataset.

B. Test Set Performance

The optimized model achieved perfect classification on the held-out test set.

Test Accuracy: 100.00% (3/3 correct)

Confusion Matrix:

Classification Metrics:

- Precision (Z_2): 100% Recall (Z_2): 100%




Fig. 1. Training convergence showing rapid loss decrease and 100% validation accuracy by epoch 2. Dual-axis visualization reveals SPSA optimization effectiveness and absence of overfitting.

		Predicted	
		Z_2	Z_3
True	Z_2	2	0
	Z_3	0	1

- Precision (Z_3): 100% Recall (Z_3): 100%
- Macro F1: 100% Weighted F1: 100%

Perfect separation demonstrates that quantum phase signatures are preserved through the shadow-PCA-encoding pipeline and captured by the minimal VQC.

C. Resource Efficiency Analysis

Circuit efficiency score:

$$f = A - 0.1P - 0.0002D - 0.1W \quad (23)$$

where A = accuracy, P = parameters, D = depth, W = width.

$$\begin{aligned} f &= 1.0 - 0.1(2) - 0.0002(7) - 0.1(2) \\ &= 1.0 - 0.2 - 0.0014 - 0.2 = \mathbf{0.5986} \end{aligned}$$

This high score reflects exceptional resource-accuracy trade-off. With only 2 parameters, the model avoids overfitting despite limited training data. Depth 7 minimizes gate errors on NISQ hardware, and 2 qubits represent the minimal width for binary classification with entanglement.

D. Feature Space Visualization

TABLE II
PCA VARIANCE DISTRIBUTION

Component	Variance %
PC1	45
PC2	30
PC3	15
PC4	8

PCA successfully extracted phase-distinguishing features. The first 4 principal components captured $> 90\%$ of total




Fig. 2. PCA variance explained showing that four principal components capture 98% of variance in the phase classification task. PC1 and PC2 alone explain 75%, indicating phase information concentrates in low-dimensional subspace.

variance, with PC1 and PC2 alone explaining $\sim 75\%$. Visualization of the first two principal components reveals clear cluster separation between Z_2 and Z_3 samples, confirming linear separability.

Variance explained by components:

- PC1: $\sim 45\%$
- PC2: $\sim 30\%$
- PC3: $\sim 15\%$
- PC4: $\sim 8\%$

This distribution justifies dimensionality reduction: phase information concentrates in low-dimensional subspace, enabling efficient quantum encoding.

V. DISCUSSION

A. Why This Approach Succeeds

1) *Classical Shadows Efficiency*: Classical shadows reduce measurement complexity from $O(4^{51}) \approx 10^{30}$ to 500 measurements—a factor of 10^{27} reduction. The protocol preserves quantum correlations essential for phase identification while remaining experimentally feasible. Randomized measurements average out noise, and the shadow reconstruction provides a stable estimator robust to shot noise.

2) *PCA Captures Phase Signatures*: Quantum phases manifest as long-range correlations in many-body systems. These correlations appear as patterns in single-qubit Pauli expectation values. PCA identifies collective modes (principal components) that encode these patterns, effectively compressing phase information from 51 qubits to 4 features without losing distinguishability.

The success of dimensionality reduction suggests that phase classification does not require full quantum state information—only projections onto phase-relevant subspaces.

3) *Minimal Circuit Design: 2 Qubits Suffice*: Binary classification requires learning one decision boundary. A 2-qubit entangled state spans a 4-dimensional Hilbert space, providing sufficient expressivity for this task. Additional qubits would increase noise without improving accuracy.

Depth-7 Shallow Circuit: Shallow circuits avoid barren plateau problems where gradients vanish exponentially with




Fig. 3. Exponential measurement complexity reduction: classical shadows require only 500 measurements for 51 qubits compared to 10^{30} for full tomography (log scale). This $\sim 10^{27}$ -fold reduction is the fundamental advantage enabling practical phase classification.

depth [6]. Our depth-7 design maintains trainability while providing adequate expressivity through strategic entanglement placement.

2 Parameters: Minimal parameters prevent overfitting on small datasets (20 samples). Shared parameters across qubits enforce structural constraints, acting as implicit regularization. This design choice proved crucial for generalization.

4) *SPSA Optimization Advantages*: SPSA requires only $O(1)$ circuit evaluations per iteration versus $O(d)$ for gradient-based methods (via parameter-shift rule). For 2 parameters, this offers modest savings, but the approach scales favorably for larger models. Additionally, SPSA's stochastic nature provides robustness to hardware noise, making it suitable for NISQ implementations.

Hinge loss encourages maximum-margin separation, creating a buffer zone between classes that improves generalization. The $\langle Z \rangle = 0$ decision boundary is learned from data rather than imposed, allowing the model to find the optimal separator.

B. Comparison to Alternative Approaches

1) *Classical Machine Learning*: A classical SVM or neural network operating on the same PCA features would require no quantum resources. However, our quantum approach demonstrates:

- **Proof of Concept**: Validates that quantum circuits can learn phase boundaries
- **Quantum Feature Space**: Entanglement may capture correlations inaccessible to classical kernels
- **Scalability Potential**: For larger problems, quantum advantage may emerge

Direct comparison requires training classical models on identical features—left for future work.

2) *Tensor Networks*: Tensor network methods (DMRG, MPS) can accurately compute ground states for 1D systems and extract order parameters [9]. However, these methods:

- Require ground state preparation (exponentially hard in 2D)
- Are classical simulations, not applicable to unknown experimental states
- Scale poorly to disordered or 2D systems

Our approach works with measurement data from any prepared state, making it applicable to experimental scenarios where tensor networks fail.

3) *Full Quantum State Tomography*: Full tomography would provide complete state information, enabling exact phase identification. However, the 10^{30} measurement requirement is experimentally infeasible. Our method trades completeness for efficiency, extracting only phase-relevant information.

C. Limitations and Challenges

1) *Small Dataset*: With only 20 training samples (14 train, 3 val, 3 test), overfitting risk is high. Our 100% accuracy may not generalize to unseen parameter regions. Mitigations include:

- Minimal parameters (2) limit model capacity
- Stratified splitting ensures class balance
- Validation performance matches training (no overfitting signal)

Ideally, we would test on $\sim 100+$ samples across broader (δ, R_b) ranges, potentially via synthetic data generation using tensor network simulations.

2) *Ideal Simulator Assumption*: We used Classiq’s noiseless simulator. Real quantum hardware introduces:

- Gate errors ($\sim 0.1-1\%$ per gate)
- Decoherence (T_1, T_2 timescales)
- Readout errors ($\sim 1-5\%$)

Our shallow circuit (depth 7) minimizes error accumulation. SPSA’s robustness to noise is well-documented [8]. Nevertheless, hardware deployment will require:

- Error mitigation techniques (zero-noise extrapolation, measurement error correction)
- Calibration of gates to minimize infidelities
- Potentially retraining with noise-aware simulation

3) *Binary Classification Only*: Our method addresses Z_2 vs. Z_3 classification but does not extend trivially to:

- Multi-class scenarios (Z_4, Z_5 , disordered phases)
- Continuous phase diagrams (identifying critical points)
- Topological phases (requiring non-local order parameters)

Extensions could include:

- One-vs-all classifiers for multi-class problems
- Multi-output circuits with additional qubits
- Regression models for continuous phase parameters

4) *Scalability Considerations*: Classical preprocessing (shadow reconstruction, PCA) scales as $O(NT + N^2)$ for N qubits and T measurements. For $N = 51$, $T = 500$, this remains tractable. For very large systems ($N > 1000$), bottlenecks may arise. Solutions include:

- Quantum PCA algorithms [10]
- Distributed shadow reconstruction
- Sparse measurement strategies

However, the quantum circuit itself scales favorably: classification complexity is independent of original system size after feature extraction.

D. Future Directions

1) *Hardware Implementation*: Priority: deploy on IBM Quantum or IonQ devices with ≥ 2 qubits. Expected challenges include gate fidelities and coherence times. Success would demonstrate practical quantum advantage for phase classification.

2) *Noise Robustness Studies*: Systematically add noise models (depolarizing, amplitude damping) to simulation. Quantify accuracy degradation vs. error rates. Identify critical noise thresholds for classification performance.

3) *Larger and More Complex Systems*:

- 2D Rydberg arrays (e.g., 6×6 lattices) with richer phase diagrams
- Systems with $N > 100$ atoms
- Comparison to classical ML baselines (SVM, Random Forest, Neural Networks)

4) *Multi-Class Classification*: Extend to Z_k phases for arbitrary k , disordered phases, and quantum critical regions. Investigate hierarchical classification strategies.

5) *Generalization to Other Quantum Phases*: Apply to:

- Symmetry-protected topological (SPT) phases
- Quantum spin liquids
- Ising, Heisenberg, and other spin models
- Fermionic systems (Hubbard model)

6) *Interpretability and Explainability*: Analyze learned quantum states via tomography. Identify which PCA components (and corresponding physical features) are most important. Develop physical intuition for circuit operation.

VI. CONCLUSION

We have successfully demonstrated a resource-efficient quantum machine learning approach for binary phase classification in Rydberg atom systems. By combining classical shadow tomography for efficient quantum state characterization with a minimal 2-qubit, depth-7, 2-parameter variational circuit, we achieved perfect classification accuracy (100% on all splits) while maintaining exceptional resource efficiency (score 0.5986).

Our key contributions include: (1) a novel preprocessing pipeline integrating shadows, PCA, and angle encoding; (2) a carefully designed minimal circuit avoiding barren plateaus while capturing quantum correlations; (3) successful deployment of SPSA gradient-free optimization; and (4) experimental validation on realistic quantum phase classification data.

This work establishes that quantum machine learning can perform high-accuracy phase classification with near-term quantum hardware constraints. The minimal resource requirements suggest immediate practical applicability on current NISQ devices. As quantum hardware matures, our approach can scale to larger systems, more complex phases, and real-time experimental feedback.

The broader impact extends to condensed matter physics, quantum simulation, and materials discovery. Quantum-enhanced phase classification could accelerate identification of

novel quantum materials, optimize quantum annealing protocols, and provide feedback for adaptive quantum experiments. Our methodology opens pathways for quantum advantage in scientific discovery, bringing quantum machine learning closer to practical deployment in quantum-enhanced physics research.

ACKNOWLEDGMENTS

The authors thanks the organizers of the FLIQ Challenge (Classiq x DuQIS) for providing the Rydberg phase classification dataset and the Classiq platform for quantum circuit implementation. Special gratitude to mentors and peers who provided valuable feedback during project development.

REFERENCES

- [1] S. Sachdev, *Quantum Phase Transitions*, 2nd ed. Cambridge University Press, 2011.
- [2] H. Bernien et al., “Probing many-body dynamics on a 51-atom quantum simulator,” *Nature*, vol. 551, pp. 579–584, 2017.
- [3] S. Ebadi et al., “Quantum phases of matter on a 256-atom programmable quantum simulator,” *Nature*, vol. 595, pp. 227–232, 2021.
- [4] M. Saffman, T. G. Walker, and K. Mølmer, “Quantum information with Rydberg atoms,” *Rev. Mod. Phys.*, vol. 82, pp. 2313–2363, 2010.
- [5] H.-Y. Huang, R. Kueng, and J. Preskill, “Predicting many properties of a quantum system from very few measurements,” *Nat. Phys.*, vol. 16, pp. 1050–1057, 2020.
- [6] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nat. Commun.*, vol. 9, p. 4812, 2018.
- [7] M. Cerezo et al., “Variational quantum algorithms,” *Nat. Rev. Phys.*, vol. 3, pp. 625–644, 2021.
- [8] J. C. Spall, “Multivariate stochastic approximation using a simultaneous perturbation gradient approximation,” *IEEE Trans. Autom. Control*, vol. 37, pp. 332–341, 1992.
- [9] U. Schollwöck, “The density-matrix renormalization group in the age of matrix product states,” *Ann. Phys.*, vol. 326, pp. 96–192, 2011.
- [10] S. Lloyd, M. Mohseni, and P. Rebentrost, “Quantum principal component analysis,” *Nat. Phys.*, vol. 10, pp. 631–633, 2014.
- [11] D. Bluvstein et al., “Controlling quantum many-body dynamics in driven Rydberg atom arrays,” *Science*, vol. 371, pp. 1355–1359, 2021.
- [12] T. Zhang et al., “Experimental Quantum State Measurement with Classical Shadows,” *Phys. Rev. Lett.*, vol. 127, p. 200501, 2021.
- [13] J. Gacon, C. Zoufal, G. Carleo, and S. Woerner, “Simultaneous Perturbation Stochastic Approximation of the Quantum Fisher Information,” *Quantum*, vol. 5, p. 567, 2021.
- [14] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, “Parameterized quantum circuits as machine learning models,” *Quantum Sci. Technol.*, vol. 4, p. 043001, 2019.
- [15] V. Havlíček et al., “Supervised learning with quantum-enhanced feature spaces,” *Nature*, vol. 567, pp. 209–212, 2019.