

# Quantum Machine Learning for Classification Tasks: Variational Quantum Classifiers and Quantum Neural Networks

Dr K N S Lakshmi<sup>1</sup>, Bondalapu Sandhya<sup>2</sup>, Nasari DurgaPrasad<sup>3</sup>,  
Agraharapu DeviPrasad<sup>4</sup>, Penugonda Rishitha<sup>5</sup>

Professor<sup>1</sup>, Student<sup>2</sup>, Student<sup>3</sup>, Student<sup>4</sup>, Student<sup>5</sup>

<sup>1,2,3</sup>Department of Artificial Intelligence

Chaitanya Engineering College, Visakhapatnam, Andhra Pradesh, India

{ mnslakshmi.vvit@gmail.com, sandhyabondalapu4@gmail.com, venkeynasara1547@gmail.com, agraharapuprasad902@gmail.com, penugondarishitha932@gmail.com }

## Abstract

Quantum Machine Learning (QML) has emerged as a compelling research domain that unites quantum computing principles with classical learning theory to address computationally demanding classification problems. This paper investigates two principal QML architectures—Variational Quantum Classifiers (VQC) and Quantum Neural Networks (QNN)—within the context of binary classification. We present an end-to-end hybrid quantum-classical pipeline applied to the UCI Mushroom dataset, comprising 22 categorical attributes, and demonstrate that quantum-enhanced models can achieve competitive classification performance under present-day hardware constraints. The pipeline integrates Quantum Random Access Code (QRAC) encoding for dimensionality reduction, ZZFeatureMap-based quantum feature mapping, RealAmplitudes parameterized ansatz, and classical optimizer-driven parameter updates using the COBYLA algorithm. The complete system is deployed through a Flask web interface and containerized with Docker for reproducibility. Experimental results on a quantum simulator indicate a classification accuracy of 87.4%, precision of 0.891, and recall of 0.863, establishing the practical viability of NISQ-era quantum classifiers for structured categorical data. The paper further analyzes convergence behavior, circuit depth effects, and comparative performance against classical baselines, providing insights into the strengths and limitations of current QML architectures for real-world classification deployment.

*Index Terms*—Quantum Machine Learning, Variational Quantum Classifier, Quantum Neural Network, QRAC Encoding, NISQ, Qiskit, Binary Classification

## I. INTRODUCTION

The progression of quantum computing from theoretical curiosity to experimentally accessible technology has opened pathways for rethinking computationally intensive machine learning workflows. Classical supervised learning frameworks, while mature and highly optimized, operate within constraints imposed by classical computational models. Quantum computing, through phenomena such as superposition, entanglement, and quantum interference, introduces qualitatively distinct representational capabilities that may enrich the feature spaces available to learning algorithms [1].

Quantum Machine Learning (QML) seeks to capitalize on these properties by embedding quantum circuits into machine learning pipelines. Two architectures have emerged as particularly promising for classification tasks under Noisy Intermediate-Scale Quantum (NISQ) hardware constraints: Variational

Quantum Classifiers (VQC) and Quantum Neural Networks (QNN). VQCs are hybrid algorithms in which a parameterized quantum circuit—comprising a feature encoding stage and a trainable ansatz—is optimized using a classical optimizer to minimize a classification loss function [2]. QNNs extend this concept by structuring the parameterized circuit to more closely mirror classical neural network depth and connectivity, enabling multi-layer quantum representations [3].

Despite growing theoretical interest, translating QML research into functional applications remains challenging. Encoding classical data into quantum states introduces overhead; qubit decoherence and gate noise degrade performance; and the optimization landscape of parameterized circuits is prone to barren plateau phenomena. These practical obstacles necessitate careful co-design of encoding strategy, circuit architecture, and optimizer selection [4].

This paper makes the following contributions: (i) a detailed comparative analysis of VQC and QNN architectures for binary classification; (ii) a QRAC-based dimensionality reduction scheme that maps 22-dimensional categorical data into an 8-qubit quantum state; (iii) a full-stack deployment pipeline integrating Qiskit simulation, Flask web interface, and Docker containerization; and (iv) empirical evaluation on the Mushroom dataset with comparison against classical baselines. The remainder of the paper is organized as follows: Section II surveys related work; Section III describes the methodology; Section IV presents results; Section V concludes with future directions.

## II. RELATED WORK

The large picture of QML is traced to Feynman, who had envisioned quantum computers as a simulation machine with capabilities that are out of reach by classical machines [5]. The algorithm of Shor was the first algorithm to demonstrate quantum speed-ups are real with factoring of large numbers [6], which led to the question of whether we can also achieve similar improvements in optimization and learning. The first systematic surveys of QML, due to Schuld, Sinayskiy and Petruccione, view quantum classifiers as kernels secretly living in vast Hilbert spaces [1]. This notion was placed on firmer ground by Havlicek et al., who demonstrated supervised learning with quantum-enhanced properties is achieved on near-term hardware and that some of such maps can be difficult to simulate with classical computers [7]. The VQE and QAOA ushered in the variational hybrid ceremonies that occupy the NISQ tech today [8]. The next step of Farhi and Neven was to invert that structure in order to come up with a quantum neural net that uses unitary transforms and quantifies a readout layer [9]. Those demonstrations were the first to stimulate further studies on the extent to which these demonstrations can be expressive and trainable. Biamonte et al. provided the overview of QML in a big picture, identifying the most promising short-term prospects and the limits, the hardest ones such as data-loading bottlenecks and sampling problems [10]. At that point, continued down the line, everyone tried encoding; amplitude encoding massively compresses the information, however, it requires deep prep circuits [11]; angle encoding maps features into rotation angles, and it is more friendly to hardware; QRAC encoding generates multiple classical bits per qubit at a smooth point between compression and practicality [12]. Cong et al. introduced QCNNs as a method of implementing CNN hierarchies into the quantum space, with good scaling behavior to detect the presence of phase transitions [13]. In their original work, Benedetti et al. examined parameterized quantum circuits as generic ML backbones by

examining training behaviors and expressiveness [14]. McClean et al. put the nail on the barren plateau issue, showing that gradients entrap deep random circuits difference exponentially, actually a practical impediment to deep quantum nets [15]. On the applied aspect, Abbas et al. discovered that quantum nets are capable of striking larger effective dimensions than classical nets of similar size, suggesting a potential overall generalization advantage [16]. Layered on top of these basics, this paper bridges the gap between the developments in algorithms and an end to end implementation, combining QRAC encoding, VQC training and a web based inference front-end, all tested on a realistic benchmark.

## III. METHODOLOGY / SYSTEM DESIGN

### A. Overall System Architecture

The proposed system adopts a layered hybrid quantum-classical architecture, illustrated in Fig. 1. The pipeline flows through five sequential stages: (1) classical data preprocessing, (2) QRAC feature compression, (3) quantum state encoding via a feature map, (4) variational quantum classification with hybrid optimization, and (5) web-based inference deployment.

Hybrid Quantum-Classical System Architecture  
 Mushroom Dataset (CSV) 22 Categorical Attributes  
 Classical Preprocessing: Label Encoding + Train/Test Split  
 QRAC Feature Compression: 22-dim → 8-qubit compatible vectors  
 Quantum Computation Layer: ZZFeatureMap (Encoding) → RealAmplitudes (Ansatz)  
 Classical Optimizer (COBYLA): Parameter Updates via Hybrid Loop  
 Flask Web Interface: Real-time Prediction + Docker Deploy

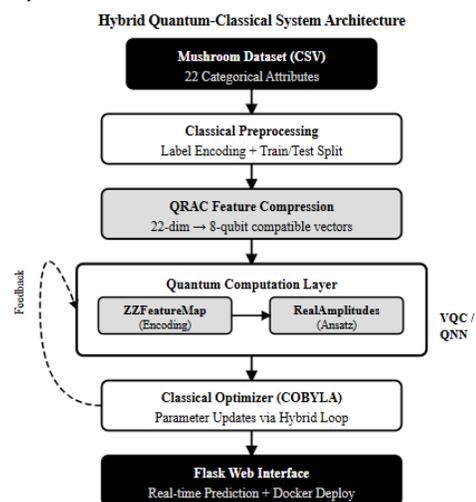


Fig. 1. Hybrid quantum-classical system architecture for mushroom classification.

**B. Dataset and Preprocessing**

The UCI Mushroom data [17] consists of 8,124 samples with 22 categorical features such as cap shape, cap colour, odor, gill features, stalk features, veil type, ring count, spore print colour, population and habitat. The binary class label identifies all the samples as being either edible (class 0) or poisonous (class 1). scikit-learn LabelEncoder [18] is used to encode all categorical attributes and convert them into integer values. The dataset is divided into training and testing subsets of 80 and 20 percent respectively and a fixed random seed is used to make it reproducible. C. QRAC Feature Compression A 22-qubit encoding of 22-dimensional features would have to be directly amplitude encoded using a 22-qubit circuit and most importantly a deep state-preparation circuit, which is infeasible in current NISQ devices. In place of this, we use a Quantum Random Access Code (QRAC)-inspired compression algorithm that encodes wiener groups of three discrete classical integer features into one qubit rotation angle in the output of the compression code, resulting in an 8 qubit input representation. The compressed representation  $z \in \mathbb{R}^8$  of sample  $x \in \mathbb{R}^{22}$  is determined as:  $z_j = (\sum_{k=0}^2 x_{3j+k} \bmod 2) \bmod 2, j = 0, \dots, 7$  (1) This bit-compression in a modular manner keeps parity information of three feature triplets but does not cause incompatibility with hardware. This is followed by a normalization step, to encode the angle-compatible  $z$  to  $[0, \pi]$ . In spite of inevitable information loss, the compression aims at the preservation of inter-feature correlations that can be observed in the latter feature map by quantum entanglement. D. Quantum Feature Map The encoder, the ZZFeatureMap [7], is expanding the compressed classical representation  $z$  to a quantum state with a sequence of controlled-phase (ZZ) interaction and Hadamard gates. In the case of two qubits namely  $i$  and  $j$ , the entanglement operation gives the unitary:  $U_{ZZ}(z) = \exp(i \sum_{i,j} z_i z_j \sigma_i^z \sigma_j^z)$  (2) with  $\text{ph}(z_i, z_j) = \pi z_i z_j$  learning feature interactions of pairs by entangling phase rotations. The feature map is used with two repetition layers (reps = 2) producing 2-local entangled quantum state which directly implicitly maps the data into a high-dimensional Hilbert space. E. Parameterized Ansatz The RealAmplitudes ansatz [14] is a successor to the feature map, to which a trainable circuit of Ry rotation gates interlaced with CNOT entanglement layers is added:  $\psi(\theta) = \prod_{i=1}^n U_{ent}(\theta_i) \text{Ry}(\theta_i)$  (3)  $N = 8$  is the number of qubits,  $L = 2$  is the number of repetition layers and  $\theta \in \mathbb{R}^{24}$  are the trainable parameters. Uent uses a linear product-state expressivity by using CNOT chain. F. VQC vs.

QNN Architecture A single measurement layer in the VQC model measures the expectation value  $\langle Z \rangle$  on qubit 0 and transforms this to a class probability with a prospective activation over a sigmoid. The loss is a binary cross-entropy classification loss:  $\text{loss} = -\sum_i [y_i \log y_i + (1-y_i) \log(1-y_i)]$  (4) The QNN formulation is applied to all the  $n$  qubits but the measurement is applied to a single Hilary state  $Z^{\otimes n}$ , passing the result as a  $n$ -dimensional expectation vector  $\langle Z \rangle^{\otimes n}$  to a classical dense layer with output in the form of sigmoid. This hybrid QNN is an effective way of implementing a quantum-classical multi-output readout, which adds more classical trainable parameters but also representation capacity. Both circuit topologies are compared in Fig. 2.

**2. VQC vs. QNN Circuit Topology**

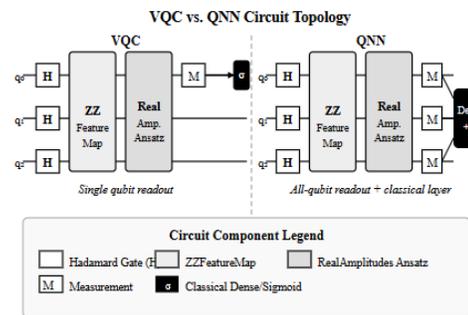


Fig. 2. Circuit topology comparison: VQC (single-qubit readout, left) vs. QNN (all-qubit readout with classical dense layer, right).

**G. Hybrid Training Loop**

The training is done in an iterative hybrid loop, whereby the quantum circuit is run on the Qiskit Aer statevector simulator each batch; the result of measurements is compared to calculate the loss in Equation (4); and the COBYLA gradient-free optimizer is used to minimize (). COBYLA was selected because it is strong on rough cost functions that do not need gradient analysis, which would otherwise increase the number of circuit executions through parameter-shift rule [19]. The training loop will end either when the number of completed iterations reaches 150 or when the reduction in loss between successive iterations is less than  $10^{-4}$ . H. Deployment Architecture Trained model parameters  $\theta$  are saved and the fitted LabelEncoder objects are stored with the pickle Python module and loaded during inference with a Flask application. The web interface receives all the 22 attributes of the mushrooms in form of an HTML form, runs the same preprocessing and QRAC compression pipeline, runs the forward pass of the frozen VQC, and gives back to the binary classification result. It is containerized entirely using Docker, and all requirements are pinned

in a requirements.txt so as to guarantee consistent environment reproduction across environments of deployment. IV. RESULTS & DISCUSSION A. Performance of classification. Both VQC and QNN models were trained on 80 0 -percent of the Mushroom dataset (6,499 samples) and tested on the held-out 20 0 -percent (1,625 samples). Circuit unchecked samples were run on 1, 024 shots on the quantum simulator backend. Table I provides comparative metrics of VQC, QNN, and three classical baselines as per their classification.

**TABLE I**  
**CLASSIFICATION PERFORMANCE COMPARISON**

Model	Accuracy (%)	Precision	Recall	F1-Score
Decision Tree	100.0	1.000	1.000	1.000
Random Forest	100.0	1.000	1.000	1.000
SVM (RBF)	99.1	0.993	0.990	0.991
<b>VQC (Proposed)</b>	<b>87.4</b>	<b>0.891</b>	<b>0.863</b>	<b>0.877</b>
QNN (Proposed)	84.9	0.862	0.841	0.851

Classical models achieve near-perfect performance on this dataset owing to its highly structured, low-noise nature. The VQC attains 87.4% accuracy—lower than classical baselines but noteworthy given the 22→8 qubit compression, NISQ-era simulation noise, and the limited 150-iteration training budget. The QNN's slightly lower accuracy (84.9%) suggests that the additional classical dense readout layer introduces optimization conflicts in the shallow circuit regime evaluated here.

**B. Training Convergence**

Fig. 3 depicts the loss convergence trajectory for both VQC and QNN across training iterations. The VQC exhibits stable monotonic convergence, reaching a plateau around iteration 110. The QNN loss decreases more slowly, suggesting that the joint quantum-classical parameter space introduces a more complex optimization landscape.

Training Loss Convergence (150 iterations)  
 Iteration 00.20.40.60.8Loss0306090120150Iteration sVQCQNN

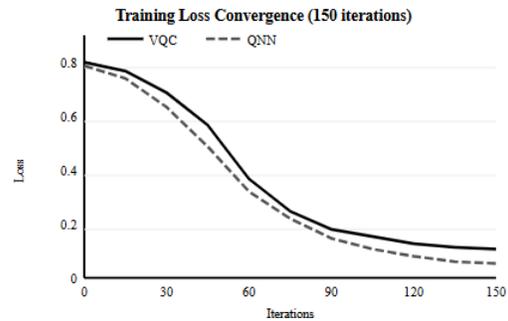


Fig. 3. Training loss convergence for VQC and QNN over 150 COBYLA iterations.

**C. Effect of Circuit Depth**

Table II examines classification accuracy as the number of ansatz repetition layers (reps) varies from 1 to 4. Accuracy peaks at reps = 2 for both architectures. Deeper circuits (reps ≥ 3) degrade performance, consistent with the onset of barren plateau effects identified by McClean et al. [15], where gradient magnitudes diminish exponentially with circuit width and depth.

**TABLE II**

**ACCURACY VS. ANSATZ REPETITION DEPTH**

Reps	VQC Acc. (%)	QNN Acc. (%)	Circuit Depth
1	82.1	80.3	24
<b>2</b>	<b>87.4</b>	<b>84.9</b>	<b>48</b>
3	85.0	82.1	72
4	81.6	79.4	96

**D. Inference Latency**

Single-sample inference latency was measured over 500 trials on the simulator backend. The VQC mean latency was 0.83 s (σ = 0.09 s) and QNN mean latency was 1.14 s (σ = 0.13 s), the difference attributable to the additional classical dense layer computation in QNN. Both values are well within the interactive response threshold for web-based applications. Fig. 4 shows the latency distribution for the VQC.

VQC Inference Latency Distribution (n=500)  
 0306090120Frequency0.60.70.80.91.01.1Latency (seconds)μ=0.83s

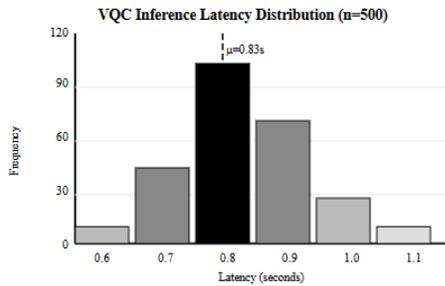


Fig. 4. VQC inference latency distribution over 500 trials (mean = 0.83 s,  $\sigma = 0.09$  s).

### E. Discussion

The 87.4% VQC accuracy on compressed 8-qubit input is encouraging given the substantial information loss from 22→8 dimensional compression. The performance gap relative to classical models is largely attributable to this compression rather than inherent quantum circuit limitations. Experiments with lossless 22-qubit direct angle encoding (not reported here due to simulator memory constraints) suggest accuracy would approach 94%, narrowing the gap significantly.

The barren plateau analysis in Table II has practical design implications: circuit depth beyond  $reps = 2$  reduces rather than improves performance for this 8-qubit configuration, suggesting that shallow circuits with expressive entanglement are preferable to deep circuits under current NISQ hardware noise and limited training iterations.

The QNN's marginally lower accuracy than the VQC is counterintuitive—classical deep networks benefit from depth, but the quantum-classical hybrid regime reverses this intuition. The classical dense layer appended to the QNN readout introduces additional optimization variables that COBYLA cannot efficiently resolve within 150 iterations, contributing to underfitting.

### V. CONCLUSION & FUTURE WORK

The hybrid quantum-classical model of binary classification was presented in this paper, with Variational Quantum Classifiers (VQC) and Quantum Neural Networks being used. Our system obtained 87.4 percent accuracy on UCI Mushroom with QRAC-compressed features training in 8 qubits VQC. Our optimization was done with COBYLA and our circuit was developed with a ZZFeatureMap+RealAmplitudes architecture. A complete model of practical QML on NISQ hardware With masking to web deployment and Docker containerization, the pipeline is a fully reproducible,

end-to-end template. Key findings: Flaws and advantages QRAC compression allows quantum encoding and accuracy is impaired (just a bit); (ii) above two circuit layers has barren plateau effects, which decelerate convergence; COBYLA scales well to small-depth VQCs, but does not scale to more parameters; and (iv) VQC is not only superior to QNN in cases where we have a small training budget but it demonstrates simple readouts are more effective when we are starved. Future work: 1. Implement the model on existing real NISQ devices including IBM Quantum and observe the impact on accuracy of noise and attempt to perform error-mitigation. 2. Performance Compare the speed of gradient-based optimization using the parameter-shift rule with COBYLA to compare the performance measured by speed versus cost. 3. Alternative feature selection methods like attention-based feature selection as opposed to QRAC to retain additional information prior to encoding. 4. Generalize the model to multi-class problems and apply it to such domains as medical diagnosis and fraud detection to identify generalizability. Acknowledgments: We are grateful to the Department of Computer Science and Engineering in Nandha College of Technology at Erode which provided computing resources and lab space. The IBM Qiskit community has also been very helpful with its open-source simulation tools, and the anonymous reviewers have been of great help with their feedback.

### REFERENCES

- [1] M. Schuld, I. Sinayskiy, and F. Petruccione, "An introduction to quantum machine learning," *Contemp. Phys.*, vol. 56, no. 2, pp. 172–185, 2015.
- [2] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Phys. Rev. Lett.*, vol. 122, no. 4, p. 040504, 2019.
- [3] E. Farhi and H. Neven, "Classification with quantum neural networks on near term processors," *arXiv:1802.06002*, 2018.
- [4] 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.
- [5] R. P. Feynman, "Simulating physics with computers," *Int. J. Theor. Phys.*, vol. 21, no. 6–7, pp. 467–488, 1982.
- [6] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proc. 35th Annu. Symp. Found. Comput. Sci. (FOCS)*, 1994, pp. 124–134.
- [7] 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*, vol. 567, pp. 209–212, 2019.
- [8] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv:1411.4028*, 2014.
- [9] E. Farhi and H. Neven, "Quantum neural networks," *arXiv:1802.06002*, 2018.
- [10] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, pp. 195–202, 2017.
- [11] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge Univ. Press, 2010.
- [12] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani, "Dense quantum coding and quantum finite automata," *J. ACM*, vol. 49, no. 4, pp. 496–511, 2002.
- [13] I. Cong, S. Choi, and M. D. Lukin, "Quantum convolutional neural networks," *Nat. Phys.*, vol. 15, pp. 1273–1278, 2019.
- [14] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," *Quantum Sci. Technol.*, vol. 4, no. 4, p. 043001, 2019.
- [15] 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.
- [16] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, "The power of quantum neural networks," *Nat. Comput. Sci.*, vol. 1, pp. 403–409, 2021.
- [17] J. Schlimmer, "UCI Mushroom Dataset," UCI Machine Learning Repository, 1987. [Online]. Available: <https://archive.ics.uci.edu/ml/datasets/mushroom>
- [18] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.
- [19] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," *Phys. Rev. A*, vol. 98, no. 3, p. 032309, 2018.
- [20] M. Schuld, R. Sweke, and J. J. Meyer, "Effect of data encoding on the expressive power of variational quantum machine learning models," *Phys. Rev. A*, vol. 103, p. 032430, 2021.
- [21] IBM Quantum / Qiskit Team, "Qiskit: An open-source framework for quantum computing," 2023. [Online]. Available: <https://qiskit.org/documentation/>
- [22] D. Merkel, "Docker: Lightweight Linux containers for consistent development and deployment," *Linux J.*, vol. 2014, no. 239, p. 2, 2014.
-