OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

Variational Quantum Classifiers: Comparing Data Encoding Strategies Under Noise

Shaurya Gera
shauryagera41@gmail.com

ABSTRACT

Variational quantum classifiers (VQCs) are promising near-term quantum machine-learning models whose performance can depend on both the data-encoding method and circuit noise. This study compared angle encoding, amplitude encoding, and data re-uploading under three levels of simulated depolarizing noise using a fixed four-qubit, three-layer VQC. The comparison was repeated across the Iris versicolor-versus-virginica and Banknote Authentication datasets using multiple stratified data splits and parameter initializations, with a classical multilayer perceptron as a baseline. The relative performance of the encoding methods depended on the dataset: angle encoding produced the highest mean accuracy on Iris, while data re-uploading produced the highest mean accuracy on Banknote Authentication. Although statistical tests identified overall encoding differences in several Banknote conditions, no pairwise comparison remained significant after multiple-comparison correction. The classical baseline outperformed all tested quantum conditions. These findings show that encoding choice can affect VQC classifier behavior under simulated noise, but they do not establish a universally superior encoding method or direct robustness of the encoded quantum states.

INTRODUCTION

Variational quantum circuits (VQCs) constitute a widespread approach to building classification models on Noisy Intermediate-Scale Quantum (NISQ) computing devices. Unlike fault-tolerant quantum algorithms that are designed to run on large-scale error-corrected quantum computers, variational quantum algorithms combine quantum state manipulation and classical parameter optimization. Thus, they can be executed on today's NISQ hardware, which still suffers from decoherence, gate errors, measurement instability, and lack of available qubits. For these reasons, examining VQC behavior in noisy conditions becomes essential to evaluate the practical relevance of near-term quantum machine learning models.

One of the central components in quantum machine learning is the data encoding task, which involves mapping classical input features onto quantum states. Unlike traditional classical machine learning, in quantum machine learning this task plays a crucial role since it establishes a way of information transfer into the quantum circuit. Data encoding is an important determinant of the quantum model's expressiveness, trainability, and sensitivity to noise. Several data encoding strategies have already been

considered in prior work, such as angle encoding, amplitude encoding, and data re-uploading. While these approaches have been analyzed previously, they were typically considered either independently of one another or without explicit noise conditions. LaRose and Coyle [9] previously investigated quantum data encodings for binary classification under both noiseless and noisy conditions. However, additional controlled evidence is useful for determining whether encoding-performance patterns remain consistent across datasets, repeated data partitions, and parameter initializations.

This study benchmarked angle encoding, amplitude encoding, and data re-uploading under three depolarizing-noise levels using a variational quantum classifier implemented in PennyLane. The experiment used two independent four-feature binary datasets: Iris versicolor versus virginica and Banknote Authentication. Each condition was evaluated using six repeated stratified data splits and six parameter initializations while the circuit architecture, optimizer, measurement procedure, and noise placement were held constant across paired comparisons. Mean test accuracy, confidence intervals, paired accuracy changes under noise, and statistical comparisons were examined. The purpose was to determine whether descriptive encoding-performance patterns were reproducible across datasets rather than to claim that one encoding was universally superior.

LITERATURE REVIEW

Variational Quantum Circuits

Variational quantum circuits are one of the key learning models for near-term quantum ML because of the interaction of quantum computation with classical optimization techniques. Variational quantum algorithms are widely described as one of the main approaches to NISQ-era computation because they rely on parameterized quantum circuits trained using classical optimizers and can therefore operate on hardware with limited qubit counts and imperfect gates [3]. Parameterized quantum circuits can also be treated as machine learning models because their adjustable rotation gates make it possible to approximate nonlinear relationships in data [2]. In VQCs, different layers of parameterized rotation gates modify the state of qubits, entanglement gates like CNOT operations provide correlations between qubits, and measurement operations transform the quantum state into a classical output. During training, the parameters of the circuit are updated using gradient optimization techniques like the parameter-shift rule.

For example, hardware-efficient variational quantum circuits have been implemented on superconducting quantum processors, demonstrating the applicability of shallow parameterized circuits on near-term devices [8]. However, trainability remains a core issue for VQCs. Prior work has shown that random initialization of variational quantum circuits can lead to the barren plateau problem, where parameter gradients become exponentially small and optimization becomes extremely difficult [10]. Together, these studies confirm that VQCs are promising yet vulnerable learning models, the efficiency of which depends on circuit architecture, parameterization, optimization, and noise resistance.

Data Encoding Methods

Data encoding refers to the process of transforming classical input features into quantum representations within a VQC. Quantum feature maps can use high-dimensional quantum state space for supervised classification, including through angle-based feature maps that generate quantum kernels [5]. This connects to the broader view of quantum machine learning as a feature-mapping process into Hilbert space, where the encoding method strongly influences the representational capacity of the learning model [12]. Data re-uploading extends this idea by repeatedly introducing classical data into the circuit, making it possible for even a single-qubit model to act as a universal quantum classifier [11]. The importance of encoding is further supported by theoretical work showing that the choice of data-encoding gates determines the Fourier frequencies available to variational quantum models [13].

Collectively, these studies suggest that the choice of data encoding significantly influences the expressibility and representational capacity of quantum learning models.

LaRose and Coyle [9] previously investigated robust quantum data encodings under noisy classification conditions, and more recent work has continued to review the practical role of quantum data-encoding methods [16]. However, additional controlled evidence is useful for determining whether encoding-performance patterns under depolarizing noise remain consistent across datasets and repeated experiments.

Noise in Quantum Machine Learning

Noise is one of the main barriers to the practical implementation of quantum machine learning because current quantum systems are affected by decoherence, imperfect gates, and measurement errors. Local Pauli noise, including depolarizing-type noise, can cause variational quantum algorithms to suffer from barren plateaus, where gradients vanish exponentially as system size increases [15]. This finding is important because it shows that noise does not only reduce final accuracy; it can also damage the optimization landscape and make training more difficult. Other work has shown that the effect of noise depends on circuit structure and optimization objective, meaning that some variational algorithms may adapt to noisy conditions better than others [14]. Noise can also change the behavior of optimized circuits, while overparameterized circuits may sometimes show higher noise resilience [4].

Altogether, these studies show that noise is not only a hardware problem but also a learning problem. Depolarizing-type noise can make training harder by flattening the optimization landscape, but not every variational circuit responds to noise in the same way [15]. Some circuits may be more resilient depending on their structure, objective function, and ability to adapt during training [4], [14]. This matters for the present study because data encoding is one of the main design choices that shapes how information enters

the circuit. Therefore, this benchmark examined whether angle encoding, amplitude encoding, and data re-uploading behaved differently when tested under the same depolarizing noise model.

Benchmarking in Quantum Machine Learning

Benchmarking is needed in quantum machine learning because proposed quantum models may perform well under simulation conditions but still fail to outperform classical counterparts. Comparisons between quantum neural networks and classical feedforward neural networks show that classical baselines are necessary when evaluating possible quantum advantages [1]. Quantum machine learning performance also depends strongly on dataset properties, since some tasks that appear difficult for classical models may become learnable when enough data are available [6]. More broadly, parameterized quantum learning models, including data re-uploading circuits, can be analyzed through a linear quantum model perspective, which helps clarify how different quantum models should be compared [7].

The above mentioned works demonstrate the need for comparison between quantum and classical approaches in terms of efficiency. However, most of the existing benchmarks include comparisons of models' architectures, optimizers, datasets, or classical approaches but do not include encoding as the main experimental factor. The present study extended this benchmarking approach by keeping the circuit size, optimizer, measurement scheme, and noise model constant while varying the encoding approach and depolarizing noise level across repeated experiments.

Gap Statement

Previous studies have investigated VQCs, data-encoding methods, quantum noise, and QML benchmarking, including work that directly examined quantum encodings under noisy conditions [9]. The remaining need is therefore not a first-ever comparison, but further controlled evidence using repeated experiments and more than one dataset. The present study addressed this narrower gap by comparing angle encoding, amplitude encoding, and data re-uploading using the same circuit architecture, optimizer, measurement procedure, and depolarizing-noise model across two four-feature datasets and six paired repeat configurations.

METHODOLOGY

Assumptions and Scope

Several design assumptions were made to maintain controlled comparisons. Both datasets were restricted to four numerical features so that the same four-qubit architecture could be used without dimensionality reduction; this improved comparability but limited generalization to higher-dimensional or multiclass tasks. A balanced sample of 100 observations per dataset was treated as sufficient for a small controlled

benchmark; matching sample size and class balance reduced dataset-size confounding, but using only a subset of Banknote and 20 test observations per repeat increased uncertainty. Six paired repeat configurations were selected as a compromise between reproducibility and the computational cost of 108 quantum-model runs; this provided repeated estimates but limited statistical power. The fixed circuit, optimizer, training budget, 60/20/20 split, and 0.5 classification threshold were applied identically across conditions so that encoding and noise level remained the principal experimental factors; conclusions may change under other modeling choices. Independent depolarizing noise was used as a simplified, reproducible approximation of hardware noise, not as a complete model of device behavior.

Four qubits were used because both datasets contained four input features. Three variational layers were selected to keep the circuit shallow and the 108 mixed-state simulations computationally manageable. The optimizer step size, training limit, validation interval, and patience were kept fixed across all quantum conditions for a controlled comparison. The noise levels represented noiseless, lower-noise, and strong-noise simulation conditions rather than the calibrated error rates of a specific device. The MLP was also kept intentionally compact as a reference baseline rather than an optimized classical model.

Datasets

Two four-feature binary classification datasets were used. The first was the Iris dataset [17], loaded through scikit-learn [19]. The setosa class was excluded, leaving versicolor and virginica as the two target classes. The resulting dataset contained 100 observations, with 50 observations from each class and four continuous features: sepal length, sepal width, petal length, and petal width. The second dataset was the Banknote Authentication dataset obtained from OpenML using data identifier 1462 [18]. During each repeat, 50 observations were sampled without replacement from each Banknote class so that the two datasets had the same balanced sample size. Six repeat configurations were evaluated using data-split seeds 42, 7, 13, 21, 99, and 123 and corresponding parameter-initialization seeds 142, 107, 113, 121, 199, and 223. For each Banknote repeat, a random number generator initialized with the corresponding data-split seed selected 50 observations without replacement from each binary class. The resulting 100 observations were combined and randomly shuffled. A stratified 20% test set was first created using the data-split seed, after which the remaining 80 observations were divided into 60 training and 20 validation observations using a second stratified split with random state equal to the data-split seed plus 10,000. The same procedure was applied to the Iris observations, except that all 50 observations from each of its two retained classes were already included. For binary classification, Iris versicolor observations were assigned label 0 and virginica observations were assigned label 1. For Banknote Authentication, the two original target labels returned by OpenML were sorted and mapped to binary labels 0 and 1, respectively. StandardScaler [19] was fitted using only the training observations and was then applied to the validation and test observations. Within each dataset and repeat, every encoding and noise condition used the same sampled observations and the same train-validation-test partition.

Circuit Architecture

The variational quantum classifier was implemented in PennyLane [20] using the default.mixed simulator. The device used four wires with shots=None, so Pauli-Z expectation values were evaluated analytically rather than estimated from a finite number of measurement shots. The QNode used PennyLane's Autograd interface with backpropagation differentiation (diff_method="backprop"). The circuit had four qubits and three variational layers. Each layer included trainable RY rotation gates, after which linear entanglement using CNOT between adjacent qubits was performed. Since there were four trainable rotation gates in each of three layers, in total the model had 12 trainable parameters. The trainable parameters were sampled from a standard normal distribution using six initialization seeds: 142, 107, 113, 121, 199, and 223. Within each repeat, the identical initial parameter matrix was used for all encoding and noise conditions so that the comparisons were paired.

Instead of measuring only one qubit, the circuit measured the expectation value of the Pauli-Z operator from all four qubits and averaged the results to produce the final model output. Let $\bar{z}$ denote the mean of the four Pauli-Z expectation values. Because $\bar{z}$ ranges from -1 to 1, it was transformed into the classifier probability $q = (\bar{z} + 1)/2$, giving q a range from 0 to 1. Binary cross-entropy was calculated using the binary target labels and this probability. Probabilities were clipped to the interval $[10^{-7}, 1 - 10^{-7}]$ before the logarithms were evaluated. For classification, $q \geq 0.5$ was assigned to class 1 and $q < 0.5$ was assigned to class 0. The model was trained using GradientDescentOptimizer with a step size of 0.1 and binary cross-entropy loss. The benchmark used full-batch training with all 60 training observations included in each optimization step. Each model was trained for a maximum of 150 steps. Validation loss was evaluated before training and at intervals of five training steps. The parameter values associated with the lowest validation loss were retained, and training stopped when the validation loss had not improved by at least 0.00001 for approximately 40 training steps. The test set was not used for optimization, early stopping, or model selection. Test accuracy was calculated only after the best validation-loss checkpoint had been selected.

Encoding Methods

Three encoding techniques were evaluated in this study: angle encoding, amplitude encoding, and data re-uploading. In angle encoding method, four standardized features were encoded into four qubits with RY rotation gates only once at the layer 0, hence there were four encoding gate operations. In amplitude encoding, the feature vector was encoded into amplitudes of the quantum state using PennyLane's AmplitudeEmbedding() function. Since four qubits require state vector of length $2^4 = 16$, the four-dimensional feature vector was zero-padded to 16 dimensions using pad_with=0.0 and then normalized to unit norm using normalize=True. In the data re-uploading method, four standardized features were encoded with RY gates at every layer instead of once. Hence, the data re-uploading method had 12 encoding gate operations, calculated as four RY encoding gates across three layers.

Noise Model

Depolarizing channel implemented in PennyLane was used as noise model for the experiments with noisy quantum hardware. As the depolarizing noise randomly perturbs the state of the qubits, it provides a good simplified representation of quantum hardware noise. In the quantum circuit, a depolarizing channel was applied to each of the four qubits separately after each variational layer. As there were three layers in the circuit, three noise injections were made for each qubit during each forward pass. Three levels of depolarization were used in the experiments: $p = 0.00$ for the noiseless case, $p = 0.01$ for the moderate noise case, and $p = 0.10$ for the high noise case. Each encoding strategy was tested at all three noise levels across two datasets and six repeated configurations, producing 108 quantum-model runs. Applying the same depolarizing channel after each variational layer provided a controlled logical-level comparison among the encoding methods, but it did not model differences in their physical hardware implementations. On a real quantum processor, amplitude-state preparation, one-time angle encoding, and repeated data re-uploading can require different gate sequences, circuit depths, and state-preparation procedures and may therefore experience different amounts or types of physical noise. Consequently, the simulated results should not be interpreted as direct hardware-level rankings of encoding noise resilience.

Classical Baseline and Evaluation Metrics

A multilayer perceptron (MLP) from scikit-learn was trained as a classical baseline for each dataset and repeated split. The MLP contained one hidden layer with eight units, used the corresponding initialization seed as its random state, and was trained for a maximum of 1,000 iterations. It received the same standardized training and test observations as the quantum models. The primary evaluation metric was test classification accuracy. For every dataset, encoding method, and noise level, the six repeated test accuracies were summarized using the mean, sample standard deviation, standard error, and 95% t confidence interval. Accuracy change under noise was calculated within each paired repeat by subtracting the noisy test accuracy from the corresponding noiseless test accuracy. These changes measured classifier behavior and were not treated as direct measurements of encoded-state robustness. Friedman tests were used to test for overall differences among the three encoding methods within each dataset and noise level. Two-sided paired Wilcoxon signed-rank tests were used for pairwise comparisons, and the three pairwise p-values within each comparison family were adjusted using the Holm procedure with the significance level set to 0.05. The statistical tests were implemented using SciPy [21]. The same statistical procedure was used to compare accuracy changes under noise.

The experiment assumed that the same four-qubit, three-layer architecture, optimizer settings, and depolarizing-noise model provided a controlled basis for comparing the three encodings. These fixed settings improved comparability but may not have been optimal for each encoding individually. The depolarizing channel was also a simplified noise model and did not reproduce all errors present on physical quantum hardware.

RESULTS

The classical MLP achieved a mean test accuracy of 0.942 ± 0.038 on the Iris versicolor–virginica dataset and 0.958 ± 0.058 on the Banknote Authentication dataset. Its mean test accuracy exceeded the mean of every quantum encoding and noise condition on both datasets.

On Iris, angle encoding produced the highest descriptive mean accuracy at each noise level: 0.750 ± 0.110 at p = 0.00, 0.733 ± 0.125 at p = 0.01, and 0.642 ± 0.124 at p = 0.10. Data re-uploading produced mean accuracies of 0.700 ± 0.089, 0.675 ± 0.099, and 0.633 ± 0.082, while amplitude encoding produced means of 0.475 ± 0.104, 0.475 ± 0.104, and 0.492 ± 0.092. The Iris Friedman tests were not statistically significant at p = 0.00, p = 0.01, or p = 0.10, with p-values of 0.054, 0.260, and 0.084, respectively. Therefore, the Iris ranking remained descriptive rather than statistically confirmed.



Figure 1. Mean test accuracy of angle encoding, amplitude encoding, and data re-uploading on the Iris versicolor–virginica dataset under three depolarizing-noise probabilities. Error bars represent 95% t confidence intervals across six paired repeat configurations.

On Banknote Authentication, data re-uploading produced the highest descriptive mean accuracy at all three noise levels: 0.867 ± 0.088, 0.858 ± 0.092, and 0.817 ± 0.129. Angle encoding produced means of 0.833 ± 0.061, 0.833 ± 0.061, and 0.717 ± 0.147, while amplitude encoding produced means of 0.592 ± 0.132, 0.600 ± 0.130, and 0.600 ± 0.100. Friedman tests identified overall differences among the Banknote encoding methods at all three noise levels, with p-values of 0.0037, 0.0037, and 0.019. However, none of the pairwise Wilcoxon comparisons remained significant after Holm correction; the smallest adjusted p-value was 0.09375. The significant omnibus tests therefore did not establish that one specific encoding was statistically superior to another.

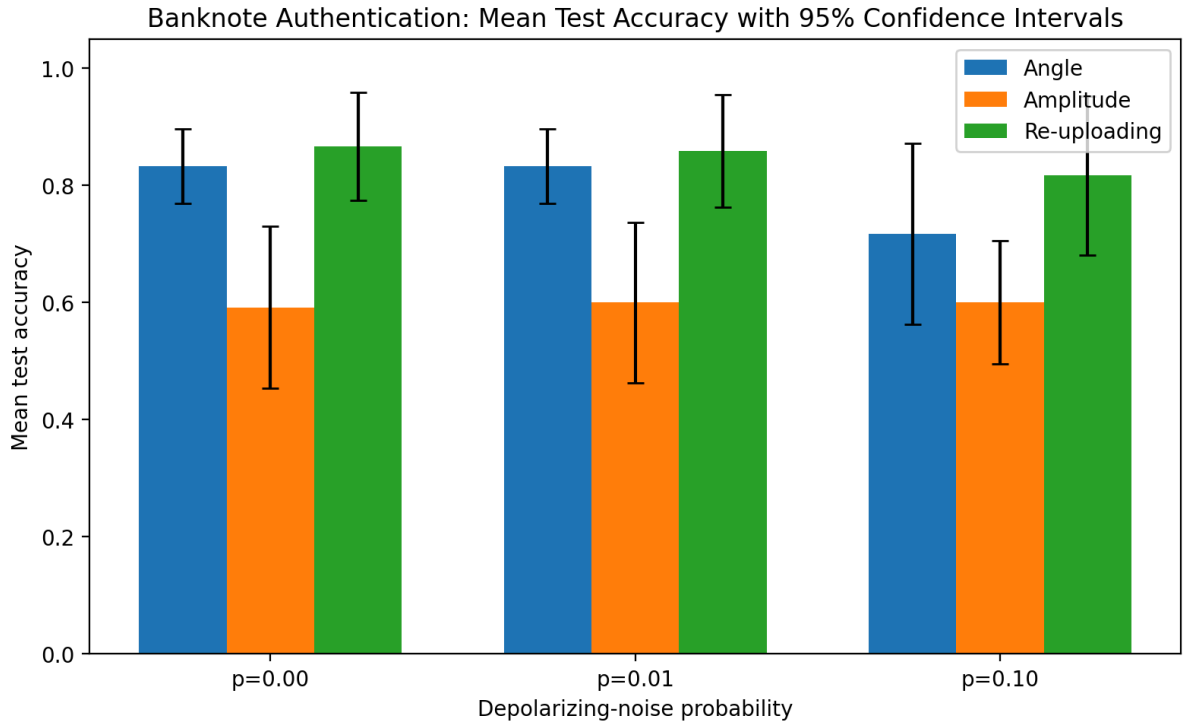


Figure 2. Mean test accuracy of angle encoding, amplitude encoding, and data re-uploading on the Banknote Authentication dataset under three depolarizing-noise probabilities. Error bars represent 95% t confidence intervals across six paired repeat configurations.

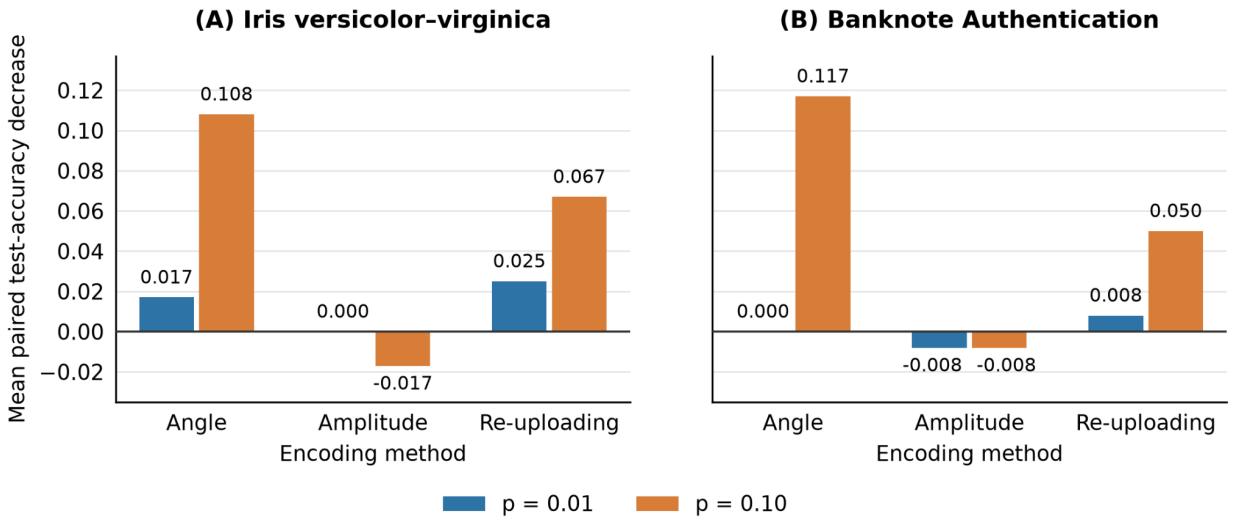


Figure 3. Mean paired decrease in test accuracy from the noiseless condition for (A) Iris versicolor–virginica and (B) Banknote Authentication. Positive values indicate lower accuracy under noise; negative values indicate higher noisy-condition mean accuracy.

Under $p = 0.10$, the Friedman test for the Iris accuracy changes produced $p = 0.042$, indicating an overall difference among the three methods. However, none of the pairwise accuracy-change comparisons remained significant after Holm correction. The other accuracy-change Friedman tests were not statistically significant. Overall, angle encoding had the highest descriptive mean accuracy on Iris, whereas data re-uploading had the highest descriptive mean accuracy on Banknote, so the results did not support one universal ranking across both datasets.

DISCUSSION

The relative performance of the encoding methods depended on the dataset. Angle encoding produced the highest mean classifier accuracy on Iris, whereas data re-uploading produced the highest mean classifier accuracy on Banknote. This change in ranking shows that the results did not identify one encoding method as universally superior. Angle encoding introduced the four input features once, while data re-uploading introduced them in each of the three layers. However, the experiment did not directly measure the effect of encoding-gate count, circuit expressibility, state fidelity, or gradient behavior. The observed accuracy differences therefore cannot be attributed causally to the number of encoding operations.

Amplitude encoding produced lower mean classifier accuracy than angle encoding and data re-uploading on both datasets. Its Iris means ranged from 0.475 to 0.492, while its Banknote means ranged from 0.592

to 0.600. Because amplitude encoding began with comparatively low noiseless accuracy, the small or negative accuracy changes under noise should not be interpreted as evidence that its encoded quantum state was better preserved. Although amplitude encoding distributes normalized feature information across quantum-state amplitudes, the experiment did not measure noisy-versus-noiseless state fidelity or another state-level quantity. Any explanation involving disruption of that distributed representation therefore remains an untested hypothesis.

The classical MLP achieved higher mean test accuracy than every tested quantum condition. On Iris, the strongest quantum mean was 0.750, compared with 0.942 for the MLP, a difference of approximately 19 percentage points. On Banknote, the strongest quantum mean was 0.867, compared with 0.958 for the MLP, a difference of approximately 9 percentage points. These results show that the tested VQC configurations did not provide a predictive advantage over the selected classical baseline. They do not establish that classical models are universally superior to quantum classifiers because only one compact VQC architecture and one MLP configuration were evaluated.

There are several limitations that should be considered when interpreting the benchmark results. First, only two binary datasets were evaluated, and both contained four numerical features, so the findings may not generalize to higher-dimensional or multiclass datasets. Second, each condition was evaluated using six paired repeat configurations. Although this improved reproducibility compared with a single run, six paired observations provided limited statistical power, especially after Holm correction. Third, each test partition contained only 20 observations, meaning that one changed prediction altered a repeat's accuracy by 0.05. Fourth, the Banknote experiment used a balanced 100-observation subset during each repeat rather than the complete dataset. Fifth, only one four-qubit, three-layer circuit architecture and one shared set of optimization settings were tested. Sixth, the experiment simulated only independent depolarizing noise and did not reproduce all errors found on physical quantum hardware or encoding-specific differences in physical state preparation, gate count, and circuit depth. Finally, the study measured classifier accuracy rather than state fidelity, trace distance, gradient behavior, or another direct mechanistic quantity. The findings therefore describe predictive behavior under the tested conditions but do not determine the physical mechanism responsible for that behavior.

CONCLUSION

This study benchmarked angle encoding, amplitude encoding, and data re-uploading under three depolarizing-noise levels using a four-qubit, three-layer variational quantum classifier. The revised benchmark included two four-feature binary datasets, six repeated stratified train-validation-test splits, six parameter initializations, validation-based model selection, 95% confidence intervals, and paired nonparametric statistical tests.

The descriptive encoding ranking depended on the dataset. Angle encoding produced the highest mean test accuracy on Iris, with values of 0.750, 0.733, and 0.642 across the three noise conditions, whereas

data re-uploading produced the highest mean test accuracy on Banknote Authentication, with values of 0.867, 0.858, and 0.817. Friedman tests identified overall encoding differences on Banknote at all three noise levels and an overall difference in high-noise Iris accuracy changes. However, no pairwise Wilcoxon comparison remained statistically significant after Holm correction. The experiment therefore did not establish that any one encoding was statistically superior across the tested configurations. The classical MLP achieved higher mean accuracy than every quantum condition on both datasets.

These findings are limited to the selected datasets, circuit architecture, optimization procedure, measurement rule, training budget, and simulated depolarizing-noise model. They describe classifier accuracy under noise rather than direct robustness of the encoded quantum states. Future work should evaluate additional datasets, more repeated trials, alternative circuit architectures, device-specific noise models, real quantum hardware, and direct state-level or gradient-based measurements

ACKNOWLEDGEMENTS

The author thanks the research mentor Shreya Savadatti for guidance on experimental design and quantum machine-learning methodology. Generative artificial intelligence tools, including ChatGPT, were used to assist with code optimization and statistical-reporting phrasing. The author executed the experiments, checked the reported outputs, made the final interpretations, and takes responsibility for the final manuscript.

REFERENCES

[1] Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., & Woerner, S. (2021). The Power of Quantum Neural Networks. *Nature Computational Science*, *1*(6), 403–409. https://doi.org/10.1038/s43588-021-00084-1

[2] Benedetti, M., Lloyd, E., Sack, S., & Fiorentini, M. (2019). Parameterized Quantum Circuits as Machine Learning Models. *Quantum Science and Technology*, *4*(4), 043001. https://doi.org/10.1088/2058-9565/ab4eb5

[3] Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., & Coles, P. J. (2021). Variational Quantum Algorithms. *Nature Reviews Physics*, *3*(9), 625–644. https://doi.org/10.1038/s42254-021-00348-9

[4] Fontana, E., Fitzpatrick, N., Ramo, D. M., Duncan, R., & Rungger, I. (2021). Evaluating the Noise Resilience of Variational Quantum Algorithms. *Physical Review A, 104(2), 022403.* https://doi.org/10.1103/physreva.104.022403

[5] Havlíček, V., Córcoles, A. D., Temme, K., Harrow, A. W., Kandala, A., Chow, J. M., & Gambetta, J. M. (2019). Supervised Learning With Quantum-enhanced Feature Spaces. *Nature*, *567*(7747), 209–212. https://doi.org/10.1038/s41586-019-0980-2

[6] Huang, H., Broughton, M., Mohseni, M., Babbush, R., Boixo, S., Neven, H., & McClean, J. R. (2021). Power of Data in Quantum Machine Learning. *Nature Communications*, *12*(1), 2631. https://doi.org/10.1038/s41467-021-22539-9

[7] Jerbi, S., Fiderer, L. J., Nautrup, H. P., Kübler, J. M., Briegel, H. J., & Dunjko, V. (2023). Quantum Machine Learning Beyond Kernel Methods. *Nature Communications*, *14*(1), 517. https://doi.org/10.1038/s41467-023-36159-y

[8] Kandala, A., Mezzacapo, A., Temme, K., Takita, M., Brink, M., Chow, J. M., & Gambetta, J. M. (2017). Hardware-efficient Variational Quantum Eigensolver for Small Molecules and Quantum Magnets. *Nature*, *549*(7671), 242–246. https://doi.org/10.1038/nature23879

[9] LaRose, R., & Coyle, B. (2020). Robust Data Encodings for Quantum Classifiers. *Physical Review A, 102(3), 032420.* https://doi.org/10.1103/physreva.102.032420

[10] McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren Plateaus in Quantum Neural Network Training Landscapes. *Nature Communications*, *9*(1), 4812. https://doi.org/10.1038/s41467-018-07090-4

[11] Pérez-Salinas, A., Cervera-Lierta, A., Gil-Fuster, E., & Latorre, J. I. (2020). Data Re-uploading for a Universal Quantum Classifier. *Quantum*, *4*, 226. https://doi.org/10.22331/q-2020-02-06-226

[12] Schuld, M., & Killoran, N. (2019). Quantum Machine Learning in Feature Hilbert Spaces. *Physical Review Letters*, *122*(4), 040504. https://doi.org/10.1103/physrevlett.122.040504

[13] Schuld, M., Sweke, R., & Meyer, J. J. (2021). Effect of Data Encoding on the Expressive Power of Variational Quantum-machine-learning Models. *Physical Review A, 103(3), 032430.* https://doi.org/10.1103/physreva.103.032430

[14] Sharma, K., Khatri, S., Cerezo, M., & Coles, P. J. (2020). Noise Resilience of Variational Quantum Compiling. *New Journal of Physics*, *22*(4), 043006. https://doi.org/10.1088/1367-2630/ab784c

[15] Wang, S., Fontana, E., Cerezo, M., Sharma, K., Sone, A., Cincio, L., & Coles, P. J. (2021). Noise-induced Barren Plateaus in Variational Quantum Algorithms. *Nature Communications*, *12*(1), 6961. https://doi.org/10.1038/s41467-021-27045-6

[16] Ranga, D., Rana, A., Prajapat, S., Kumar, P., Kumar, K., & Vasilakos, A. V. (2024). Quantum Machine Learning: Exploring the Role of Data Encoding Techniques, Challenges, and Future Directions. Mathematics, 12(21), 3318. https://doi.org/10.3390/math12213318

[17] Fisher, R. A. (1936). Iris [Dataset]. UCI Machine Learning Repository. doi:10.24432/C56C76.

[18] OpenML. Banknote Authentication, OpenML dataset 1462.

[19] Pedregosa, F., Varoquaux, G., Gramfort, A., et al. (2011). Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research, 12, 2825–2830.

[20] Bergholm, V., Izaac, J., Schuld, M., et al. (2018). PennyLane: Automatic differentiation of hybrid quantum-classical computations. arXiv:1811.04968.

[21] Virtanen, P., Gommers, R., Oliphant, T. E., et al. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods, 17, 261–272. doi:10.1038/s41592-019-0686-2.