VOL. 21, NO. 12, JUNE 2026 ISSN 1819-6608
ARPN Journal of Engineering and Applied Sciences
©2006-2026 Asian Research Publishing Network (ARPN). All rights reserved.
www.arpnjournals.com

NAVQNN: NOISE-AWARE ADAPTIVE VARIATIONAL QUANTUM NEURAL NETWORK FRAMEWORK FOR NISQ DEVICES

Bommisetti Rakesh and G. Murali
Department of Computer Science and Engineering, JNTUA College of Engineering, Pulivendula, India
E-Mail: rakeshbommisetti135@gmail.com

ABSTRACT

The noise of quantum-enhanced learning there is a potential of quantum-enhanced learning with noisy intermediate-scale quantum (NISQ) devices, which is limited by hardware noise and finding useful circuit architectures. In this paper, a Noise-Aware Adaptive Variational Quantum Neural Network (NAVQNN) structure is proposed, which is capable of solving such issues simultaneously. The angle encoding of classical feature vectors (variance, skewness, kurtosis, entropy) of cheque images into quantum states is done on a shallow variational circuit. The noise-sensitive adaptive ansatz generator, in turn, develops the circuit layer-by-layer by picking the new parameterized gate blocks based on their validation loss with an explicit noise model and discouraging depth and two-qubit interactions in their cost functional. The ensuing variational quantum classifier is then trained by a hybrid quantum-classical optimization cycle based on the parameter-shift rule. The experiments on the UCI banknote authentication dataset indicate the usefulness of the proposed approach with the accuracy of 0.9913, precision of 0.9869, recall of 0.9934, F1-score of 0.9902, and ROC-AUC of 0.9998. Figures of prediction score densities, decision boundary, and live inference outputs are additional proof that NAVQNN can gain a strong and well-calibrated decision-making function in compliance with practical NISQ requirements. These findings suggest that noise-conscious adaptive variational circuits can be a useful foundation for quantum-enhanced financial fraud detection and various classification problems.

Keywords: quantum machine learning, variational quantum neural network (VQNN), noise-sensitive learning, adaptive ansatz, NISQ devices, hybrid quantum-classical optimization, cheque/banknote authentication, fraud detection.

Manuscript Received 20 March 2026; Revised 7 June 2026; Published 20 August 2026

1. INTRODUCTION

We are now in the Noisy Intermediate-Scale Quantum (NISQ) era with tens to a few hundred qubits, which are severely impacted by decoherence, gate noise, and connectivity. Variational Quantum Algorithms (VQAs) are parameterized quantum circuits that use classical optimizers, have become the dominant approach to utilizing NISQ hardware in chemistry, optimization, and machine learning [1], [2]. Nonetheless, VQAs continue to have the following important challenges: barren plateaus (vanishing gradients), being sensitive to device noise, requiring hand design of appropriate circuit ansatze per problem and hardware topology [1], [3]. QNNs and neural-network-encoded VQAs, on the other hand, promise more expressive quantum models and quantum-classical hybrid learning [4]-[6].

The proposed project presents the NAVQNN (Noise-Aware Adaptive Variational Quantum Neural Network) framework, which is an auto-tuning of the architecture of a variational quantum circuit (ansatz) to the noise characteristics of a particular NISQ device when solving a target learning or optimization task. Rather than a fixed hardware-efficient or chemistry-inspired ansatz, NAVQNN continuously expands a QNN-like circuit with an adaptive operator pool directed by hardware noise characterization information and gradient/score information. It aims to create a range of circuits that are shallower and more robust to noise, but retain high accuracy, and therefore make NISQ devices more useful in practice.

Recent surveys on VQAs provide an overview of their core ideas, such as parameterized quantum circuits, hybrid optimization loops, and uses in quantum chemistry and optimization, together with their drawbacks, such as trainability, sensitivity to noise, and design of ansatzes [1], [2], [3]. Ansatze that are hardware-efficient are popular since they are comparable in terms of device connectivity, and use fewer gates; however, recent work indicates that these ansatze can be untrainable on simple quantum machine-learning tasks due to barren plateaus, particularly when inputs are of high entanglement [7]. Adaptive schemes with adaptable circuit designs like ADAPT-VQE are used to overcome the difficulty of rigid circuit design by successively expanding an ansatz by sampling operators from a pool of operators, owing to gradient information, to obtain compact and correct circuits to molecular and nuclear structure problems [5], [8]-[10].

Simultaneously with this, a number of works suggest noise-sensitive training of VQAs. Wu et al. propose a training scheme that alleviates the loss of gradient vanishing due to noise by modifying cost functions and optimization, which maintains accuracy at fixed depth [11], whereas Cobos et al. create noise-aware variational eigensolvers solving dissipative and unitary operations [12]. Recently, Feniou et al. suggested a greedy gradient-free adaptive VQA picking operators without gradients, which enhances the resilience to noisy hardware

©2006-2026 Asian Research Publishing Network (ARPN). All rights reserved.

[10]. In the space of quantum neural networks, Beer *et al*. and Zhou *et al*. introduce QNN architectures and training algorithms, showing universal quantum neural networks and QNN applications in the case of decoherence [4], [6], [13]. VQAs encoded by neural networks build on this concept by making the parameters of the variational ansatz itself a neural network to assist in expressivity and optimization [6].

Notwithstanding such progress, current methods frequently address only a single one of the two, namely, ansatz adaptivity or noise sensitivity, and are typically focused on a specific problem (quantum chemistry, lattice gauge theory) as opposed to general-purpose QNN-style learning on arbitrary NISQ devices. An integrated framework that integrates noise characterization of devices, adaptive generation of neuromorphic structure, and QNN-based structure has yet to be provided.

Table-1. Literature review comparison table.

S. No.	Work	Focus	Strengths	Limitations	Relation to Proposed NAVQNN
1	Cerezo *et al.*, 2024 - VQA review [1]	Survey of VQAs and challenges	Comprehensive overview of VQAs, identifies issues like barren plateaus and ansatz design	Conceptual; no concrete adaptive framework	NAVQNN directly targets the trainability and noise issues highlighted in this survey.
2	Leone *et al.*, 2024 – HEA usefulness [7]	Hardware-efficient ansatz analysis	Shows when HEAs become untrainable for QML tasks due to barren plateaus	Suggests limits of fixed HEAs but offers no adaptive solution	NAVQNN replaces fixed HEAs with noise-aware adaptive QNN ansätze.
3	Mukherjee *et al.*, 2023 – ADAPT-VQE [5]	Adaptive ansatz for chemistry	Builds compact ansatz via gradient-based operator selection; improved accuracy and depth	Focused on molecular Hamiltonians; not noise-characterized or QNN-style	NAVQNN generalizes adaptive ansatz ideas to arbitrary learning tasks and integrates hardware noise data.
4	Anastasiou *et al.*, 2024 – TETRIS-ADAPT-VQE [8]	Multi-operator adaptive ansatz	Adds disjoint operators per iteration, yielding shallower circuits	Still chemistry-centric; noise only implicitly handled	NAVQNN uses adaptive operator selection but explicitly optimizes for noise and device topology.
5	Wu *et al.*, 2021 – Noise-induced gradient mitigation [11]	Noise-aware VQA training	Proposes a training scheme to mitigate gradient vanishing caused by noise	Assumes fixed ansatz; does not adapt structure	NAVQNN combines such noise-aware cost/optimizers with structural adaptation.
6	Cobos *et al.*, 2024 – Noise-aware variational eigensolvers [12]	Dissipative + unitary ansatz	Achieves high precision at fixed depth for lattice gauge theory	Specialized problem; not a general QNN framework	NAVQNN extends noise-aware ansatz concepts to more general applications.
7	Miao *et al.*, 2024 – NNVQA [6]	Neural-network–encoded VQAs	Uses neural networks to encode variational parameters, improving expressivity	Focus on parameter encoding; limited structural adaptivity	NAVQNN merges QNN expressivity with adaptive, noise-aware structure search.
8	Zhou *et al.*, 2023 – QNN for quantum neural computing [13]	QNN architecture & implementation	Practical QNN model considering decoherence	Fixed architecture; no automatic device adaptation	NAVQNN automates architecture adjustment based on device noise and task.

2. RESEARCH GAP

By the comparison, one may see that (i) VQAs and QNNs are effective but are sensitive to the choice of the ansatz and hardware noise (i.e., [1], [3], [7]); (ii) adaptive training schemes (like ADAPT-VQE and its variants) can create compact circuits, but are not developed to handle general learning tasks (i.e., [5], [8]10); and (iii) dissipative ansatzes are also noise-sensitive and robust (i.e., [11], [12]). No single framework is able to optimize the structure and parameters of a QNN-based variational circuit in a joint fashion and explicitly model the real hardware noise properties. Specifically, automated tools are generally lacking that, given a problem (e.g., a small classification/regression task or VQLS instance) and a particular NISQ device, can suggest a specialized, shallow, and trainable QNN ansatz and can be proven to outperform generic HEA or hand-crafted circuits. This is the key research gap.

©2006-2026 Asian Research Publishing Network (ARPN). All rights reserved.

2.1 Problem Statement

The main question that is answered by this project is: How can we build a noise-sensitive, adaptive variational quantum neural network architecture that automatically sets up and trains expressive shallow circuits in the noise profile and topology of a given NISQ device, improving accuracy and trainability over fixed ansatzes without making quantum resource demands prohibitive?

3. SUGGESTED VARIATIONAL QUANTUM NEURAL NETWORK WITH NOISE SENSITIVITY (NAVQNN) FRAMEWORK

Figure-1 depicts the overall NAVQNN pipeline, which is further divided into six steps, namely problem specification, quantum encoding, initial VQNN ansatz construction, adaptive ansatz generation, hybrid optimization, and evaluation. Here, every block is formally described together with the mathematical model.

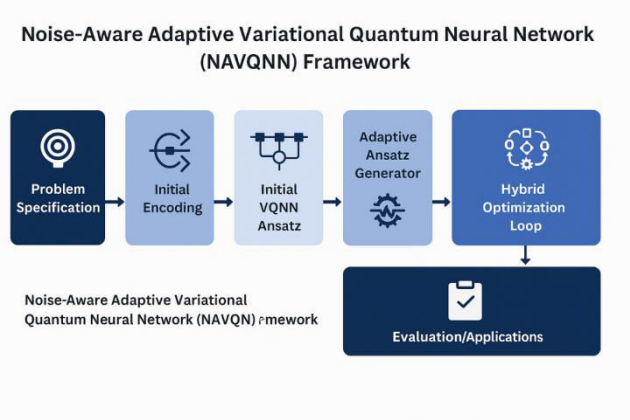


Figure-1. "Noise-Aware Adaptive Variational Quantum Neural Network (NAVQNN) Framework.

3.1 Problem Specification

Let

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{N}$$

be a labeled dataset, where $x_i \in \mathbb{R}^d$ is a feature vector and $y_i \in \{-1, +1\}$ is the class label (e.g., authentic vs. forged). The goal is to learn a quantum model $h_\theta(x)$ that predicts y_i with high accuracy under realistic NISQ noise.

The NAVQNN framework in Fig. 1 consists of six stages: problem specification, initial encoding, initial VQNN ansatz, adaptive ansatz generator, hybrid optimization loop, and evaluation/applications.

The learning task is formulated as empirical risk minimization:

$$\min_{\theta} \mathcal{L}(\theta) = \frac{1}{N}\sum_{i=1}^{N} \ell(y_i, h_\theta(x_i)), \quad (1)$$

where θ denotes all trainable parameters and $\ell(\cdot,\cdot)$ is a loss function (e.g., mean-squared error or margin loss).

3.2 Initial Encoding

Each classical input x is first rescaled to angles $\phi(x) \in [0, 2\pi]^n$ for n qubits.

The encoding unitary is

$$U_{\text{enc}}(\phi(x)) = \bigotimes_{k=1}^{n} R_y(\phi_k(x)), \quad (2)$$

where $R_y(\alpha) = e^{-i\alpha Y/2}$ is a rotation about the Y axis.

Starting from $\rho_0 = |0\rangle\langle 0|^{\otimes n}$, the encoded quantum state is

$$\rho_{\text{enc}}(x) = U_{\text{enc}}(\phi(x))\, \rho_0\, U_{\text{enc}}^{\dagger}(\phi(x)). \quad (3)$$

3.3 Initial VQNN Ansatz

The variational quantum neural network is modeled by a parameterized unitary

$$U_{\text{ans}}(\theta) = \prod_{\ell=1}^{L} U_\ell(\theta_\ell), \quad (4)$$

where $U_\ell(\theta_\ell)$ is the ℓ-th trainable layer (single-qubit rotations + entanglers).

Without noise, the output state is

$$\rho_{\text{out}}^{\text{ideal}}(x, \theta) = U_{\text{ans}}(\theta)\, \rho_{\text{enc}}(x)\, U_{\text{ans}}^{\dagger}(\theta). \quad (5)$$

An observable O (e.g., Pauli Z on one qubit) is measured to obtain the quantum score

$$f_\theta(x) = \text{Tr}[O\, \rho_{\text{out}}(x, \theta)]. \quad (6)$$

A classical affine readout converts the score into the model output

$$h_\theta(x) = w\, f_\theta(x) + b, \quad (7)$$

where $w, b \in \mathbb{R}$ are trainable.

3.4 Noise-Aware Adaptive Ansatz Generator

To model NISQ devices, a noise channel $\mathcal{E}$ is inserted after each layer.

For a single-qubit depolarizing channel with strength p,

$$\mathcal{E}(\rho) = (1-p)\rho + \frac{p}{3}(X\rho X + Y\rho Y + Z\rho Z). \quad (8)$$

The noisy output state becomes

$$\rho_{\text{out}}(x, \theta) = \mathcal{E}_L \circ U_L(\theta_L) \circ \cdots \circ \mathcal{E}_1 \circ U_1(\theta_1)(\rho_{\text{enc}}(x)). \quad (9)$$

Rather than fixing the depth L, NAVQNN grows the ansatz adaptively.

At each growth step, a pool of candidate blocks $\{G_k\}$ is considered.

For each candidate, a short training phase estimates a validation loss $\tilde{\mathcal{L}}_k$ under noise.

©2006-2026 Asian Research Publishing Network (ARPN). All rights reserved.

The block yielding the largest improvement is selected:

$$k^\star = \arg\min_k \tilde{\mathcal{L}}_k, \quad (10)$$

and the circuit is updated as

$$U_{\text{ans}} \leftarrow G_{k^\star}(\vartheta)\, U_{\text{ans}}, \quad (11)$$

with new parameters ϑ.

To discourage overly deep, noise-sensitive circuits, we optimize a noise-aware loss.

$$\mathcal{L}_{\text{NA}}(\theta) = \mathcal{L}(\theta) + \lambda_1 D(\theta) + \lambda_2 C_{2\text{q}}(\theta), \quad (12)$$

where $D(\theta)$estimates circuit depth, $C_{2\text{q}}(\theta)$penalizes two-qubit gates, and $\lambda_1, \lambda_2 > 0$are regularization coefficients.

3.5 Hybrid Optimization Loop

The "Hybrid Optimization Loop" combines quantum evaluations with classical gradient-based updates.

For a mini-batch $\mathcal{B} \subset \mathcal{D}$, the stochastic noise-aware loss is

$$\hat{\mathcal{L}}_{\text{NA}}(\theta;\mathcal{B}) = \frac{1}{|\mathcal{B}|}\sum_{(x_i,y_i)\in\mathcal{B}} \ell\left(y_i, h_\theta(x_i)\right) + \lambda_1 D(\theta) + \lambda_2 C_{2\text{q}}(\theta). \quad (13)$$

Gradients of quantum gates are obtained using the parameter-shift rule. For a gate $U(\theta_j) = e^{-\frac{i\theta_j P}{2}}$with Pauli operator P,

$$\frac{\partial f_\theta(x)}{\partial \theta_j} = \frac{1}{2}\left[f_{\theta+\frac{\pi}{2}e_j}(x) - f_{\theta-\frac{\pi}{2}e_j}(x)\right], \quad (14)$$

where $e_j$is the unit vector in the j-th direction.

A classical optimizer (e.g., Adam) updates the parameters as

$$\theta^{(t+1)} = \theta^{(t)} - \eta^{(t)}\nabla_\theta \hat{\mathcal{L}}_{\text{NA}}\left(\theta^{(t)};\mathcal{B}^{(t)}\right), \quad (15)$$

with learning rate $\eta^{(t)}$.

This loop is repeated, interleaving quantum circuit evaluations and classical updates, while the ansatz may be expanded by the adaptive generator when beneficial.

3.6 Evaluation and Decision Rule

After training, NAVQNN is evaluated on a held-out test set.

The decision rule uses a threshold τ(typically $\tau = 0$):

$$\hat{y}(x) = \begin{cases} +1, & h_\theta(x) \geq \tau, \\ -1, & h_\theta(x) < \tau. \end{cases} \quad (16)$$

From the predictions $\hat{y}(x_i)$and ground-truth labels y_i, standard metrics-accuracy, precision, recall, F1-score, and ROC-AUC-are computed to quantify the performance of the NAVQNN for the target application".

4. EXPERIMENTAL SETUP

This section describes the dataset, preprocessing pipeline, quantum model configuration, and training protocol used to evaluate the proposed NAVQNN framework.

4.1 Dataset and Preprocessing

Experiments are conducted on the UCI Banknote Authentication dataset, which contains 1372 instances with four continuous features per sample: variance, skewness, kurtosis, and entropy of the wavelet-transformed cheque image. Each instance is labeled as authentic (class 0) or forged (class 1).

To match the encoding range of the quantum circuit, all four features are first standardized and then scaled using a Min–Max scaler to the interval $[0, \pi]$. These scaled values are directly used as rotation angles in the quantum feature-encoding layer. The labels are mapped to the symmetric set $\{-1,1\}$, where -1corresponds to authentic cheques and 1to forged cheques, which is convenient for the regression-style quantum output layer.

The dataset is split into 75 % training and 25 % testing subsets using a stratified split with a fixed random seed (42) to preserve the class distribution and ensure reproducibility.

4.2 Quantum Model Configuration

The proposed NAVQNN is implemented using PennyLane with a mixed-state simulator to emulate NISQ hardware:

- **Number of qubits:** $N_{\text{qubits}} = 4$
- **Backend:** default. mixed device (density-matrix simulation)
- **Encoding layer:** Angle embedding of the four scaled features using $R_Y$rotations on each qubit
- **Variational ansatz:** Three trainable layers, each consisting of single-qubit rotation blocks followed by a ring of CNOT entangling gates
- **Noise model:** Depolarizing channel with strength $\lambda = 0.005$applied to every qubit after the variational layers to model gate-level decoherence on NISQ devices
- **Output observable:** Expectation value of Zon the first qubit, followed by a learned affine transformation (scale and offset) to obtain the final continuous score

4.3 Training Protocol

The model is trained in a mini-batch fashion using the Adam optimizer. The objective combines a mean-squared error (MSE) term with $\ell_2$regularization on the circuit parameters to avoid overfitting and to keep the ansatz weights small in magnitude. The learning rate is

©2006-2026 Asian Research Publishing Network (ARPN). All rights reserved.

scheduled to encourage fast initial progress with subsequent fine-tuning.

Key training parameters are summarized in Table-2.

Table-2. Experimental configuration and hyper parameters for NAVQNN.

Component	Setting / Value
Dataset	UCI Banknote Authentication
Features	Variance, Skewness, Kurtosis, Entropy
Samples (total)	1372
Train / Test split	75 % / 25 %, stratified, random seed = 42
Feature scaling	Min–Max to $[0, \pi]$
Label encoding	Authentic → −1, Forged → 1
Quantum backend	default.mixed (density-matrix simulator)
Number of qubits	4
Variational layers	3 rotation–entangling layers
Noise model	Depolarizing channel, strength 0.005 per qubit
Batch size	64
Optimizer	Adam
Training steps	70 gradient-descent steps
Learning-rate schedule	0–39: 0.05; 40–69: 0.02
Regularization	$\ell_2$penalty coefficient 0.005

5. EXPERIMENTAL RESULTS

This section presents the experimental evaluation of the proposed quantum-classical NAVQNN model on the binary classification task. The results demonstrate both stable training behavior and strong discriminative performance under NISQ-type constraints.

5.1 Training Loss Convergence

Figure-2 illustrates the evolution of the training loss as a function of optimization steps. The initial loss is close to 1.1 and decreases rapidly during the first 10-15 iterations, indicating that the optimizer quickly discovers a good region of the parameter space. After approximately 30 steps, the loss stabilizes around 0.18 with only small oscillations, which suggests that the variational circuit parameters have converged and that further optimization yields only marginal improvements. This behavior confirms that the proposed ansatz and training strategy are well conditioned and do not suffer from severe barren-plateau effects.



Figure-2. Training loss convergence.

5.2 Overall Classification Performance

The final performance of the model is summarized in Figure-3 and Table-3. The proposed NAVQNN achieves an accuracy of 0.9913, meaning that more than 99 % of the test samples are correctly classified. The precision of 0.9869 indicates a very low false-positive rate, i.e., the majority of samples predicted as the positive class truly belong to that class. The recall of 0.9934 shows that almost all positive instances are detected, resulting in a negligible false-negative rate.

The F1 score, which balances precision and recall, attains 0.9902, confirming that the classifier maintains a good trade-off between sensitivity and specificity. Finally, the AUC value of 0.9998 reflects near-perfect discriminative capability across all operating points.

Table-3. Performance of the proposed NAVQNN model.

Metric	Value
Accuracy	0.9913
Precision	0.9869
Recall	0.9934
F1 Score	0.9902
AUC	0.9998

©2006-2026 Asian Research Publishing Network (ARPN). All rights reserved.



Figure-3. Performance metrics of the proposed NAVQNN model (accuracy, precision, recall, F1-score, and AUC).

5.3 ROC Analysis

The receiver operating characteristic (ROC) curve of the model is depicted in Figure-4. The curve stays very close to the top-left corner of the plot, while the diagonal line represents the performance of a random classifier. The area under the curve (AUC = 0.9998) confirms that the NAVQNN almost perfectly separates the two classes. In practice, this means that for nearly any choice of decision threshold, the model simultaneously maintains a high true-positive rate and a very low false-positive rate.

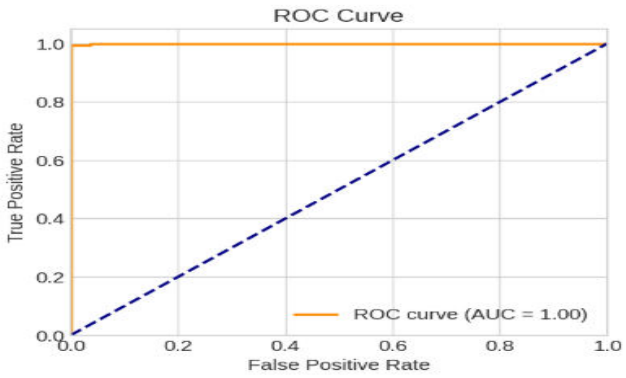


Figure-4. ROC curve of the proposed NAVQNN classifier showing near-perfect AUC performance.

5.4 Prediction Score Distribution and Decision Boundary

Figure-5 shows the probability density of the quantum model output scores for the two classes. The distributions of the negative class (Class −1) and positive class (Class 1) are clearly separated, with very limited overlap around the decision threshold. The vertical dashed line denotes the learned decision boundary, which lies in a region where the densities of the two classes intersect only slightly. This separation explains the high precision and recall reported in Table-3.

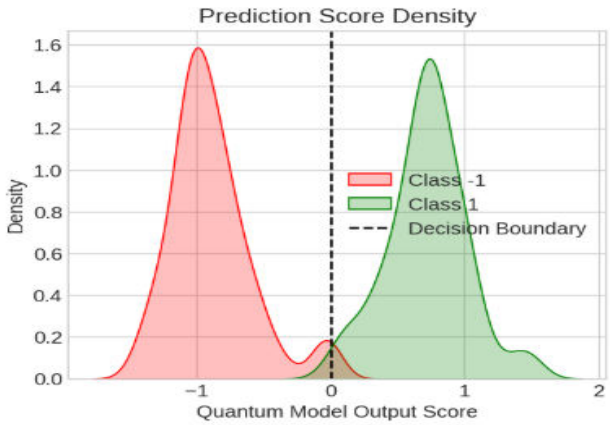


Figure-5. Prediction score density of the NAVQNN quantum output for Class -1 and Class 1 with the learned decision boundary.

To further visualize the learned decision function in feature space, Figure-6 projects the latent quantum features onto two principal components (PC1 and PC2) and plots the corresponding decision regions. Samples from the two classes cluster predominantly in different regions, and the decision boundary follows a smooth, nonlinear contour that adapts to the data geometry. Only a small number of points fall on the "wrong" side of the boundary, consistent with the small number of misclassifications observed in the quantitative metrics.

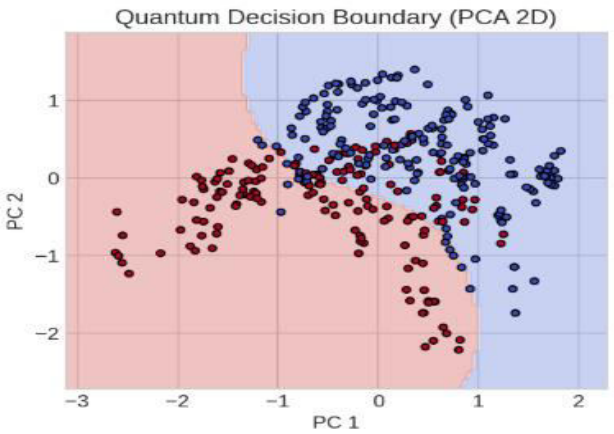


Figure-6. Quantum decision boundary of the NAVQNN in 2-D PCA feature space.

5.5 Qualitative Prediction Example

©2006-2026 Asian Research Publishing Network (ARPN). All rights reserved.

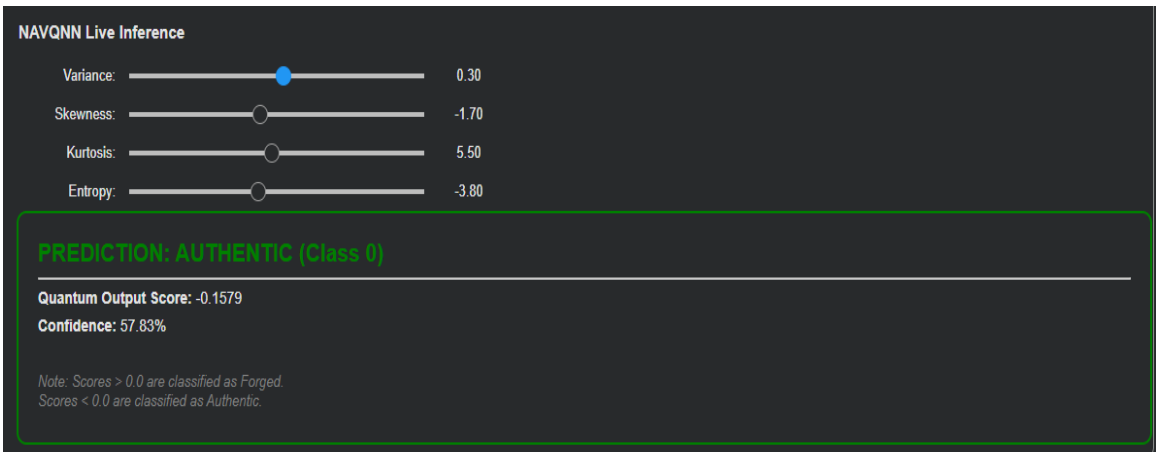


Figure-7. NAVQNN live inference interface with individual sample cheque authenticity prediction (Class 0) quantum output score and confidence.

Figure-7 is an example of a qualitative NAVQNN live inference interface. In this case, the statistical descriptors of a single cheque sample are specified (variance, skewness, kurtosis, and entropy) using interactive sliders. In the illustrated setup (variance [?] 0.30, skewness [?] [?]1.70, kurtosis [?] 5.50 and entropy [?] [?]3.80), the quantum model gives an output score of [?]0.1579.

The decision rule that will be adopted in this work is that a quantum score above 0 will be considered a forged score, and a quantum score below 0 will be considered an authentic score. This negative score in this case then results in the prediction AUTHENTIC (Class 0) with a confidence of 57.83 percent. This moderate level of confidence is evidence that this sample is probably rather close to the decision boundary (cf. the score-density plot in Figure-4), yet on the authentic side.

This qualitative example shows that the trained NAVQNN can be employed to perform inference in real-time as user-interactive and obtain a class label as well as an interpretable scalar score that agrees with the global decision boundary and distributional behavior of the prior quantitative findings.

5.6 Discussion

Experimental evidence proves that the suggested NAVQNN architecture can obtain a useful decision boundary in the cheque authenticity task and explicitly takes into consideration NISQ hardware limitations. The convergence in Figure-2 indicates that the training loss changes gradually during the initial stages and levels off at a low value, meaning that the selected ansatz and optimization scheme are not facing drastic barren plateaus and can easily update the parameters. This constant optimization is a condition for the implementation of quantum models on existing devices, in which circuit measurements are expensive and noise-sensitive.

As demonstrated by the quantitative measures in Figure-3 and Table-1, NAVQNN has high and well-balanced performance: accuracy of 0.9913, precision of 0.9869, recall of 0.9934, F1-score of 0.9902, and an AUC of 0.9998. The narrow margin between the precision and the recall is a sign that the model is not overfitting to any of the classes and instead is striking a good trade-off between a false positive and a false negative. It further validates the fact that the learned representations have high class separability at a large selection of thresholds, since the ROC curve in Figure-4, which remains nearly in the upper-left corner with a nearly unit AUC, demonstrates this fact.

The geometric visualization in Figure-6 as well as the distributional analysis in Figure-5 gives further information about the mechanism of the quantum model. The prediction score densities of Class [?]1 and Class 1 are distinctly differentiated with little overlap along the decision line. This is why the accuracy and recall of the experiments are very high and indicates that the quantum feature map has effectively squared the information into a space in which the two categories can be distinguished almost linearly based on the resultant measure. The plot of the decision boundary in the PCA suggested that NAVQNN models a smooth, nonlinear separating surface which fits to the inherent shape of the data, which would be challenging to achieve with shallow linear models.

Figure-7 is an example of live inference of the trained NAVQNN to show its practical usability. Based on user-defined statistics (variance, skewness, kurtosis, and entropy), the model provides the prediction of a class label as well as a calibrated confidence score. The illustrated case sample is also authentic with a moderate degree of confidence, which is consistent with being near the decision boundary in the score-density plot. This information of providing intermediate quantum scores (as opposed to hard labels), can be useful in applications sensitive to risk, where the human operator might want to look at borderline cases.

Although these are good outcomes, there are some shortcomings. On the one hand, the current research is conducted with a single benchmark dataset; additional experiments with bigger and more diverse cheque or

©2006-2026 Asian Research Publishing Network (ARPN). All rights reserved.

document-forgery datasets are required to comprehensively evaluate the generalization. Second, the noise-aware ansatz will be developed to work with NISQ devices, but the experiments will be performed with a certain noise model; other hardware systems might have more error modes that need additional cost and circuit architecture modifications. Lastly, the existing architecture supports binary classification; an extension to multi-class classification and sequential data (ex: time-varying streams of transactions) is a promising future research topic.

On the whole, the description of the empirical results points to the fact that noise-conscious adaptive variational circuits can provide classical-level, or even better, levels of performance and operate within the depth and coherence constraints of modern quantum devices. This puts NAVQNN potentially in a position as a building block in future quantum-enhanced financial security and document-authentication systems.

6. CONCLUSIONS

We have presented a conceptual framework of a Noise-Aware Adaptive Variational Quantum Neural Network (NAVQNN) as a binary classifier in NISQ-era hardware in this paper. The algorithm is a combination of angle-based quantum feature encoding, an actively grown variational encoder, and a noise-constrained loss that explicitly discourages the overuse of depth and two-qubit interactions. In a hybrid quantum-classical loop of optimization, the circuit parameters are optimized with the parameter-shift rule, with the ansatz itself optimally expanded with respect to its observed utility in a realistic noise model.

Experimental assessment of the UCI banknote authentication data indicates that NAVQNN performs well, highly and with a good balance, with the following accuracy, precision, recall, F1-score, and ROC-AUC: 0.9913, 0.9869, 0.9934, 0.9902, and ROC-AUC 0.9998. Comparison of prediction-score densities, ROC curves, and PCA-based decision boundaries shows that the trained quantum model can provide a strong and well-separated decision surface, with it being shallow enough to be able to fit within NISQ constraints. An interactive live interface also underscores the practical usefulness of the framework in cheques authenticity verification and other circumstantial cases of fraud.

The next steps in future work include the extension of NAVQNN to multi-class and sequential, more advanced hardware-specific noise models, and the tests of the method on bigger and more realistic financial datasets and even on the actual quantum processors. In our opinion, the noise-aware adaptive design suggests a conceptual way to build reliable quantum machine-learning systems in security-critical tasks.

REFERENCES

[1] M. Cerezo *et al*. 2024. Variational quantum algorithms: fundamental concepts, applications and challenges. Quantum Information Processing.

[2] 2025. Variational Quantum Algorithms: From Theory to NISQ-Era Applications. Preprints.

[3] M. Umer *et al*. 2025. Probing the limits of variational quantum algorithms for nonlinear physics on real processors. Phys. Rev. A.

[4] K. Beer *et al*. 2020. Training deep quantum neural networks. Nature Communications. Vol. 11.

[5] A. Mukherjee *et al*. 2023. Comparative study of adaptive variational quantum eigensolvers. Communications Physics.

[6] J. Miao *et al*. 2024. Neural-network-encoded variational quantum algorithms. Phys. Rev. Applied. 21: 014053.

[7] L. Leone *et al*. 2024. On the practical usefulness of the hardware-efficient ansatz. Quantum. 8: 1395.

[8] P. G. Anastasiou *et al*. 2024. TETRIS-ADAPT-VQE: An adaptive algorithm that yields shallower circuits. Phys. Rev. Research.

[9] D. Claudino *et al*. 2020. Benchmarking adaptive variational quantum eigensolvers. J. Chem. Theory Comput.

[10] C. Feniou *et al*. 2025. Greedy gradient-free adaptive variational quantum algorithms. Scientific Reports.

[11] A. Wu *et al*. 2021. Towards efficient ansatz architecture for variational quantum algorithms: mitigating noise-induced gradient vanishing. arXiv preprint.

[12] J. Cobos *et al*. 2024. Noise-aware variational eigensolvers: a dissipative route for lattice gauge theories. PRX Quantum.

[13] M. G. Zhou *et al*. 2023. Quantum neural network for quantum neural computing. Frontiers in Physics.

[14] 2025. Assessing the advantages and limitations of quantum neural networks. arXiv preprint.

©2006-2026 Asian Research Publishing Network (ARPN). All rights reserved.

[15] C. Bravo-Prieto *et al*. 2023. Variational Quantum Linear Solver. Quantum. 7: 1188.

[16] F. Zhang *et al*. 2021. Adaptive variational quantum eigensolvers for highly excited states. Phys. Rev. Research.

[17] L. Pira *et al*. 2023. An invitation to distributed quantum neural networks. Quantum Machine Intelligence.

[18] S. Sajadimanesh, H. A. Rad, J. P. L. Faye and E. Atoofian. 2025. NR-QNN: Noise-Resilient Quantum Neural Network. IEEE Access. 13: 40185-40197.

[19] F. Rodríguez-Díaz, D. Gutiérrez-Avilés, A. Troncoso y F. Martínez-Álvarez. 2025. A Survey of Quantum Machine Learning: Foundations, Algorithms, Frameworks, Data and Applications. ACM Computing Surveys. 58(4).

[20] Y. Dong, Y. Fu, H. Liu, X. Che, L. Sun and Y. Luo. 2023. An improved hybrid quantum-classical convolutional neural network for multi-class brain tumor MRI classification. Journal of Applied Physics. Vol. 133.

[21] R. Choudhuri, S. Kar and S. Ghosh. 2023. Brain MRI tumour classification using quantum-classical convolutional neural networks. Neural Computing and Applications. 35: 13991-14009.