Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Hybrid Quantum Neural Networks with Feature Attribution for Explainable Risk Estimation of Chronic Kidney Disease

Emmanuel Ismaila Watila

ei.watila@gmail.com

University of Maiduguri https://orcid.org/0009-0009-3445-1620

Peter Yusuf Dibal

University of Maiduguri https://orcid.org/0000-0001-9278-3994

Stephen Bassi Joseph

University of Maiduguri https://orcid.org/0000-0001-5050-0132

Research Article

Keywords: Quantum machine learning, quantum neural networks, chronic kidney disease, disease prediction, model interpretability, feature attribution

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11844357/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Hybrid Quantum Neural Networks with Feature Attribution for Explainable Risk Estimation of Chronic Kidney Disease

E.I Watila[1*], P.Y Dibal[1] and S.B Joseph[1]

[1]Department of Computer Engineering, University of Maiduguri, Maiduguri, Borno State, Nigeria.

*Corresponding author(s). E-mail(s): ei.watila@gmail.com;

Abstract

Chronic kidney disease (CKD) is a progressive, long-term condition characterized by a gradual loss of renal function. It represents a significant global health burden, particularly among elderly populations and specific ethnic groups, including individuals of Black and South Asian descent. Early detection and management are critical to preventing severe complications, yet timely and accurate diagnosis remains a clinical challenge. This study investigates the application of Quantum Machine Learning (QML) for the predictive diagnosis of CKD. The study proposes a hybrid quantum neural network (HQNN) architecture and evaluates its performance against alternative QML configurations and conventional classical machine learning techniques. Employing a benchmark CKD dataset from the University of California, Irvine (UCI) Machine Learning Repository, we implemented a robust preprocessing pipeline comprising data imputation, advanced hybrid feature selection, and data normalization. The proposed hybrid model was trained and evaluated on simulated 5-qubit noisy quantum hardware via `FakeManilaV2`. It achieved a classification accuracy of approximately 97.5% on the UCI CKD cross-sectional cohort with a high circuit expressivity in the quantum layer. The model performed at baseline for classical models and outperformed some standalone QML techniques evaluated in this study.The findings illustrate the feasibility of integrating shallow hybrid quantum-classical networks for medical diagnostic tasks under current NISQ constraints. The proposed hybrid framework provides an exploratory demonstration of explainable risk estimation for CKD with potential relevance for future clinical decision-support research once validated on larger and more diverse cohorts.

Keywords: Quantum machine learning, quantum neural networks, chronic kidney disease, disease prediction, model interpretability, feature attribution

1 Introduction

1.1 Background

The global burden of chronic kidney disease (CKD) is rising sharply in both absolute and relative terms, driving up worldwide morbidity and mortality rates. While this public health crisis severely impacts regions such as Latin America, the Caribbean, North Africa, the Middle East, and Oceania, it poses an escalating challenge for low- and middle-income nations, particularly within sub-Saharan Africa [1]. In Nigeria, however, a critical shortage of population-based data obscures the true scale of the disease, hindering the development of effective, targeted prevention and management frameworks [2].

Advanced machine learning, specifically deep learning using neural networks, has demonstrated profound utility across the medical domain. Physicians [3] and scientists [5] have routinely leveraged these computational frameworks to engineer novel solutions that facilitate highly precise clinical diagnoses and inferences. Notable applications include the automated quantification of interstitial fibrosis and tubular atrophy from kidney ultrasonography images, as well as the early detection of liver fibrosis utilizing graph convolutional networks.

1.2 Motivation

However, as classical computing approaches its physical and theoretical boundaries, the scientific and technological communities are compelled to engineer alternative computational paradigms [6],[7]. Quantum computing has emerged as a revolutionary alternative, uniquely capable of solving complex problems that remain intractable for classical hardware [8]. Even within the constraints of the current Noisy Intermediate-Scale Quantum (NISQ) era, these processors offer significant advantages. This intersection of quantum hardware and artificial intelligence has catalyzed the emergence of Quantum Machine Learning (QML), a novel discipline that synthesizes core principles from both fields to deliver substantial improvements in computational speed and algorithmic performance [9].

1.3 State of the Art

While classical machine learning has achieved high diagnostic accuracy in nephrology; deep learning approaches often suffer from overfitting when applied to small, heterogeneous, or high-dimensional clinical datasets. Quantum Machine Learning (QML) offers theoretical advantages by mapping features into high-dimensional quantum Hilbert spaces to capture non-linear relationships.

Recent literature has increasingly explored this potential for nephrological applications. For instance, an initial benchmark framework evaluated Quantum Support Vector Machines (QSVM) against classical Support Vector Machines counterparts using standard datasets applying Principal Component Analysis (PCA) and Singular Value Decomposition [10] while other approaches investigated architectures like Variational Quantum Classifiers (VQCs) and VQCs with sparse autoencoders [12], using standalone quantum-enhanced support vector machines (SVMs) for chronic kidney

disease prediction, quantum kNN, & quantum Naive-Bayes [13], [14]. However, their real-world clinical deployment remains hindered by Noisy Intermediate-Scale Quantum (NISQ) constraints, including hardware decoherence, limited qubit connectivity, and data-encoding overheads.

Furthermore, these existing methodologies exhibit critical shortcomings or gaps. Many of these architectures treat QML models as pure black boxes, leaving them severely unexplainable for real-world clinical deployment; because medical diagnosis requires explainability i.e. attributing what pushed the patient towards a diagnosis, a single binary classification doesn't provide the explainability needed for real-world clinical adoption.

To address these gaps, our contribution introduces a hybrid quantum-classical predictive framework engineered specifically to address NISQ-era limitations or constraints while preserving clear diagnostic functionality. Unlike prior models, we utilize an optimized shallow-depth ansatz and feature encoding strategy to minimize circuit depth and overhead. We validated this framework on clinical data [20] establishing a transitionally viable path to CKD risk profiling.

1.4 Contributions

In this paper we investigate the use of a hybrid quantum neural network (HQNN) for binary classification [11] for the predictive diagnosis of Chronic kidney disease (CKD). Using the University of California, Irvine Machine Learning Repository CKD dataset, we evaluate the performance of our hybrid quantum-classical approach against other standalone quantum methods like QSVM, VQCs, QkNN [10], [12], [13], [14], and classical techniques [15], [16], [17], [19].

Our primary contribution towards this proposed hybrid quantum-classical prediction framework are as follows;

1. **Design of a NISQ-Resilient Hybrid Architecture:** We developed a hardware-efficient hybrid quantum-classical predictive framework optimized specifically for current NISQ-era constraints. By employing a shallow-depth ansatz, the architecture achieves minimal gate counts and execution latency without sacrificing classification fidelity.
2. **Optimized Low-overhead quantum data-encoding:** Our methodology introduces a localized custom feature-mapping protocol that cuts down data-encoding overhead; compressing clinical features into compact quantum gate rotations rather than deep qubit chains, the model preserves vital non-linear feature interactions within the Hilbert space while maintaining a shallow circuit depth.
3. **Probabilistic Risk Estimation & Attribution Layer:** To transform model outputs into actionable clinical intelligence, our pipeline generates dual-class percentage risk scores P(CKD) and P(NOTCKD) via a softmax probability transform. When coupled with our Integrated Gradients attribution layer, clinicians receive both a calibrated risk percentage and a ranked breakdown of the exact biomarkers driving that score.

1.5 Highlight of Method

We propose a NISQ-resilient hybrid quantum-classical architecture that pairs a low-depth, 3-qubit variational circuit with targeted clinical feature selection to achieve high-accuracy CKD prediction under hardware constraints. To ensure clinical utility, the framework integrates a dual-class softmax risk estimator with an Integrated Gradients attribution layer, translating quantum outputs into calibrated risk percentages and biomarker-level explanations.

2 Related Works & Literature Review

Using variational quantum circuits as neurons, achieved 10% higher accuracy and 20% better cost minimization on simulated quantum hardware compared to individual circuits. However, performance on actual quantum hardware was limited by small qubit and gate constraints, indicating scalability issues with current technology [11].

The variational quantum classifier (VQC) achieved a higher f1-score, recall, precision, specificity, and AUC-ROC compared to other methods like support vector machine (SVM), random forest, and multilayer perceptron – with higher accuracy of 99.2% (after feature selection) for identification and classification of CKD [12].

Utilized quantum computing principle, by using a quantum support vector machine (QSVM) which achieved an accuracy of 97.50% compared to its classical counterpart, the classical support vector machine (SVM) for diagnosis of CKD [13].

Utilizing an ensemble of various quantum machine learning techniques like quantum support vector machines (QSVMs), quantum k-nearest neighbors, and quantum Naive-Bayes for detection of various chronic diseases including kidney disease achieved an overall classification accuracy of 96.8% outperforming the classical counterparts [14].

Use of machine learning techniques for detection of chronic kidney disease, such as the Logistic Regression, Random Forest, Gradient Boosting algorithms and Support Vector Machine (SVM) classifier. The final output predicts whether the person has CKD or not by using a minimum number of features by performing feature reduction using several statistical tests (ANOVA, Pearson's Correlation, & Cramer's V). The gradient boosting classifier achieved an accuracy of 99.1% [19].

Employed experiential analysis of ML techniques for classifying the kidney patient dataset as CKD or NOTCKD. Seven ML techniques together with NBTree, J48, Support Vector Machine, Logistic Regression, Multi-layer Perceptron, Naïve Bayes, and Composite Hypercube on Iterated Random Projection (CHIRP) are utilized. The experimental results using accuracy revealed 95.75% for NB, 96.50% for LR, 97.25% for MLP, 97.75% for J48, 98.25% for SVM, 98.75% for NBTree, and 99.75% for CHIRP [16].

Six machine learning algorithms (logistic regression, random forest, support vector machine, k-nearest neighbor, naive Bayes classifier, and feed-forward neural network) were used to establish models. Among these machine learning models, random forest achieved the best performance with 99.75% diagnosis accuracy [17].

Developed and validated a deep learning algorithm (DLA) to screen for chronic kidney disease non-invasively from retinal images, using data from a large population-based study, and externally validated the DLA in two independent datasets in Singapore and China. The DLA showed good performance with an AUC of 0·911 in internal validation and 0·733 and 0·835 in external test sets [18].

Data reduction used to ensure the removal of redundant features in the dataset while maintaining a high analytical result with machine learning algorithms. Feature selection and feature association was carried out. Several statistical tests were utilized, such as the ANOVA test, Pearson's correlation, and Cramer's V test [19] were used to determine the relationships between variables in the dataset.

Through this review, we found strengths and architectural limitations within the currently implemented classical and quantum machine learning techniques for clinical diagnosis & prediction. Current standalone quantum-enhanced models like QSVMs remain plagued by hardware decoherence, qubit connectivity bottlenecks, and cumulative gate errors on NISQ devices. To overcome these, we proposed adopting a hybrid quantum-classical approach that utilizes a dense classical layer for non-linear dimensionality reduction thereby reducing the number of qubits, keeping circuit depth shallow with minimal repetition on features map and ansatz; resulting in theoretical improvements in execution time within the physical T_1 and T_2 coherence windows, preventing multi-qubit entangling gates errors, and decay of quantum states into thermal noise.

3 Materials & Methodology

The benchmark dataset is a publicly available dataset on the University of California, Irvine Machine Learning Repository with an ID `336`. The dataset had a total of 400 instances — 250 ckd and 150 notckd, with 24 features and 1 class [20]. This benchmark dataset features had missing values with the highest missing rate at 38% for rbc, this necessitated the requirement for imputation. Furthermore other standard preprocessing steps like feature association and feature selection were required to select the best features to achieve a high analytic result, other preprocessing steps like scaling were also applied.

Table 1 Features present in the CKD dataset by Rubini, L, et al [20].

Feature	Detail	Type	Missingness
`age`	Age	Numerical	2.25%
`bp`	Blood Pressure	Numerical	3.00%
`sg`	Specific Gravity	Nominal	11.75%
`al`	Albumin	Nominal	11.50%
`su`	Sugar	Nominal	12.25%
`rbc`	Red Blood Cells	Nominal	38.00%
`pc`	Pus Cell	Nominal	16.25%
`pcc`	Pus Cell Clumps	Nominal	1.00%
`ba`	Bacteria	Nominal	1.00%
`bgr`	Blood Glucose Random	Numerical	11.00%
`bu`	Blood Urea	Numerical	4.75%
`sc`	Serum Creatinine	Numerical	4.25%
`sod`	Sodium	Numerical	21.75%
`pot`	Potassium	Numerical	21.75%
`hemo`	Hemoglobin	Numerical	13.00%
`pcv`	Packed Cell Volume	Numerical	17.75%
`wbcc`	White Blood Cell Count	Numerical	26.50%
`rbcc`	Red Blood Cell Count	Numerical	32.75%
`htn`	Hypertension	Nominal	0.50%
`dm`	Diabetes Mellitus	Nominal	0.50%
`cad`	Coronary Artery Disease	Nominal	0.50%
`appet`	Appetite	Nominal	0.25%
`pe`	Pedal Edema	Nominal	0.25%
`ane`	Anemia	Nominal	0.25%

3.1 Imputation

There are several ways to handle missing values in pattern mining and data-analysis algorithms, especially in the medical domain. Single Conditional Mean Imputation, Multiple Imputation, Maximum Likelihood, and Pattern-Mixture Model are additionally usable [21]; we settled on using Multiple Imputation [22].

We leveraged the Multiple Imputation with Chained Equation (MICE) on numerical & nominal values (after encoding the nominal values into numerical values) by using sklearn's `IterativeImputer` with `GradientBoostingRegressor` estimator in the scikit-learn Python library.

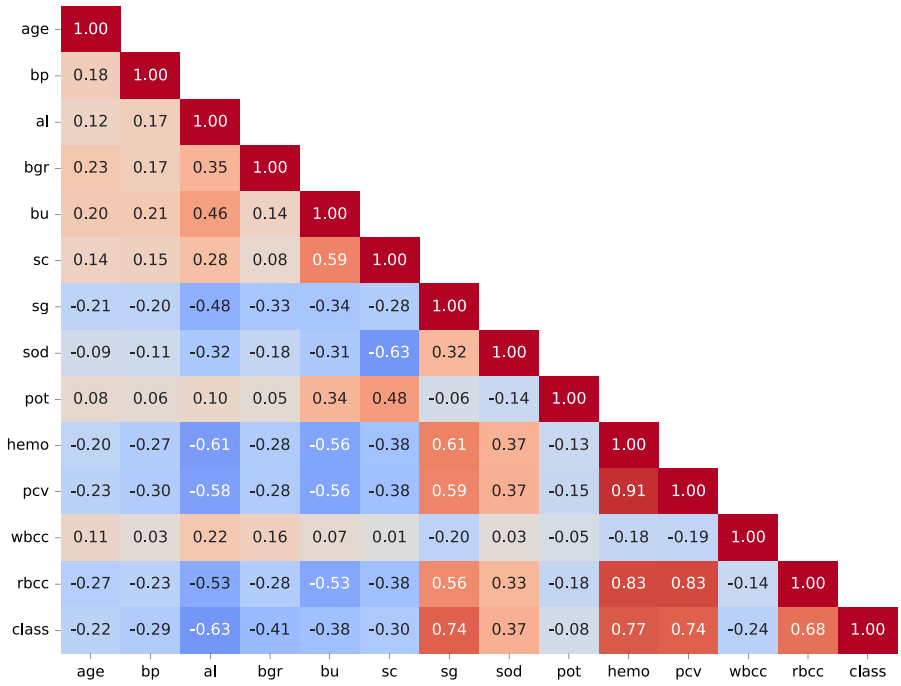

Fig. 1 Numerical Features Correlation Matrix using Pearson's Correlation

3.2 Feature Selection, Normalization & Preprocessing

The dataset had an unbalanced class, where the number of patients without CKD (notckd 150) was significantly less than the number of patients with CKD (ckd 250). The next set of concerns with this dataset was selecting relevant features and other preprocessing tasks, like normalizing the selected feature samples.

Feature selection was an essential step in reducing the input variables to the model by using only relevant data (using the least amount of predictors) and eliminating noise while preserving physiologically coupled signals; this was achieved using a hybrid correlation-based feature selection [23], using the Pearson Correlation Coefficient & Cramer's V Test [19] and domain-guided clinical prioritization on feature with strong internal collinearity [30], [31]. With strict values on both the strong positive and negative correlation ranges, we determined the threshold of ($r >= 0.5$) and ($r <= -0.4$).

Pearson Correlation Coefficient & Cramer's V was applied to explore the internal structure of the dataset and identify the relationship between each feature of the dataset and the target class. The results of the correlation tests showed the features in the dataset (numerical and nominal) that had a strong negative correlation or a strong positive correlation with the target class.

The result provided 10 features after the feature selection process, leveraging the Pearson's Correlation & Cramer's V Test. Some features fell below the feature selection criteria but were still retained in the final dataset because they exhibited strong correlation within themselves, i.e. some of the features selected through the feature

association tests exhibited typical statistical collinearity the likes of haemoglobin (hemo) and packed cell volume (pcv) with high correlation coefficient ($r = 0.91$) and serum creatinine (sc) and blood urea (bu) with $r = 0.59$ which are two features that are important biomedical & haematological markers for estimating the eGFR and BUN:Cr Ratio for staging [24], [25], [26]. Rather than discarding any of these diagnostic signals, our framework relies on a trainable dense classical layer, compressing internal collinearities into a lower-dimensional latent space to prevent model instability while capturing subtle pathophysiological discrepancies such as cell volume versus hemoglobin density.

Table 2 Selected Numerical Features using Pearson Correlation

Feature	Correlation	Remark
`bgr`	−0.415 (−0.42)	Strong Negative
`al`	−0.625 (−0.63)	Strong Negative
`hemo`	0.766 (0.77)	Strong Positive
`sg`	0.739 (0.74)	Strong Positive
`pcv`	0.739 (0.74)	Strong Positive
`rbcc`	0.667 (0.67)	Strong Positive

Table 3 Selected Nominal/Categorical Features using Cramer's V

Feature	Association	Remark
`rbc`	0.418 (0.42)	Strong Negative
`htn`	0.582 (0.60)	Strong Negative
`pc`	0.395 (0.40)	Strong Positive
`dm`	0.555 (0.60)	Strong Positive

We applied a simple transformation to represent all features on the same scale. For this, we squeezed all features onto the interval (0, 1) using `MinMaxScaler` and then passed this through a feature map [11].

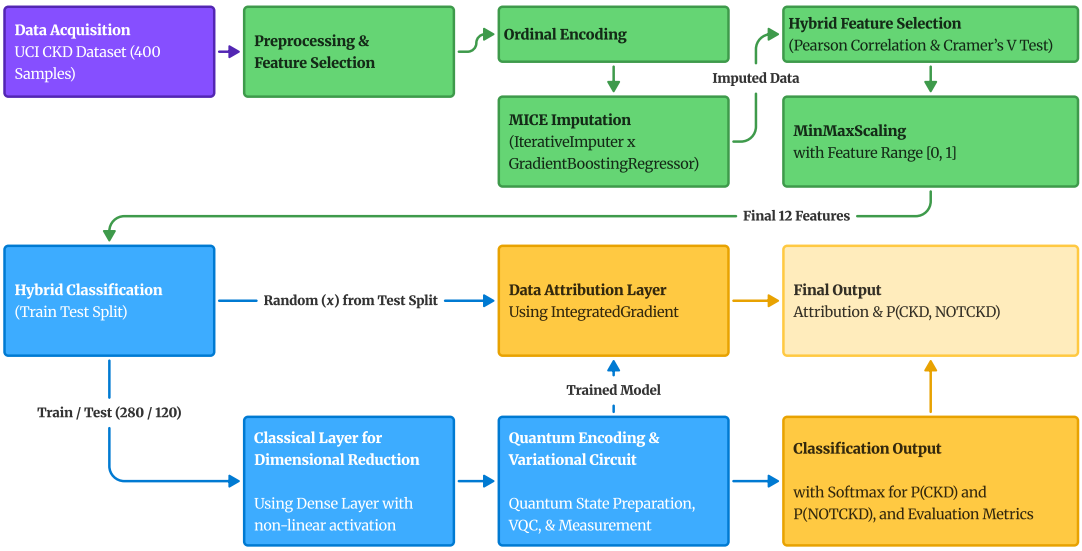

Fig. 2 Experiment Flow Diagram

3.3 Model Architecture & Quantum Circuit Design

In this section, we discuss the proposed model and its circuit design, and the physical principle governing its operation.

3.3.1 Hybrid Quantum-Classical Architecture

A Multi-Layer Perceptron (MLP) is a feed-forward artificial neural network where layers sequentially transform input features via linear operations (weights and biases) and non-linear activation functions. In our hybrid architecture, the model functions as a classical-quantum neural network, where specific classical layers are replaced by a Parameterized Quantum Circuit (PQC) acting as a variational quantum layer.

The classical-quantum interface is operationalized using Qiskit's TorchConnector wrapper, enabling seamless integration of the EstimatorQNN within PyTorch's automatic differentiation (autograd) engine. To minimize quantum circuit depth and noise on NISQ-era simulators, we utilized a 3-qubit variational circuit. The quantum state pipeline consists of classical data encoding via the ZZFeatureMap concatenated with a trainable RealAmplitudes ansatz.

3.3.2 Quantum Data Encoding using ZZFeatureMap

The ZZFeatureMap is a parameterized quantum circuit (in this study using 3-qubit) designed to map classical feature vectors into a high-dimensional quantum Hilbert space ($\mathcal{H}$). The feature map encodes data through single-qubit phase rotations (RZ) proportional to input feature values, followed by two-qubit entangling gates (CZ or $CNOT$ gates executing ZZ interactions) between pairs of qubits. These ZZ interactions capture non-linear correlations and feature interdependencies in quantum space. For this architecture, a single repetition layer ($reps = 1$) was utilized to preserve shallow circuit depth.

3.3.3 RealAmplitude (Variational Ansatz)

The RealAmplitudes circuit serves as the parameterized variational ansatz, acting as the quantum analogue to trainable weights in a classical neural network layer. It consists of alternating layers of parameterized single-qubit RY rotation gates—which constrain state amplitudes to real numbers—and entangling $CNOT$ gates across the 3 qubits. During training, backpropagation via TorchConnector updates the circuit's variational parameters alongside classical weights to minimize the loss function. A single repetition was applied to maintain expressivity and limit barren plateaus and gate noise.

3.3.4 Probabilistic Risk Estimation & Attribution Layer

The hybrid quantum-classical model was operationalized within a Flask server. We added a probabilistic risk estimation which uses `Softmax` to process the neural network's raw logits into P(CKD) and P(NOTCKD) establishing diagnostic confidence. The attribution layer uses `IntegratedGradients` to evaluate feature attribution against unnormalized model logits conditioned on the target class; this layer quantifies the precise contribution of each biomarker (features) toward the final classification and returns the integral convergence delta to guarantee attribution stability.

4 Results

The proposed hybrid quantum neural network was run on noisy hardware calibration data using `FakeManilaV2`; a 5-qubit NISQ processor and compared with other standalone & classical models. The experiments were run multiple times and the mean of the metrics were taken. The Qiskit libraries & Jupyter Notebook were used to visualize the results of the experiment.

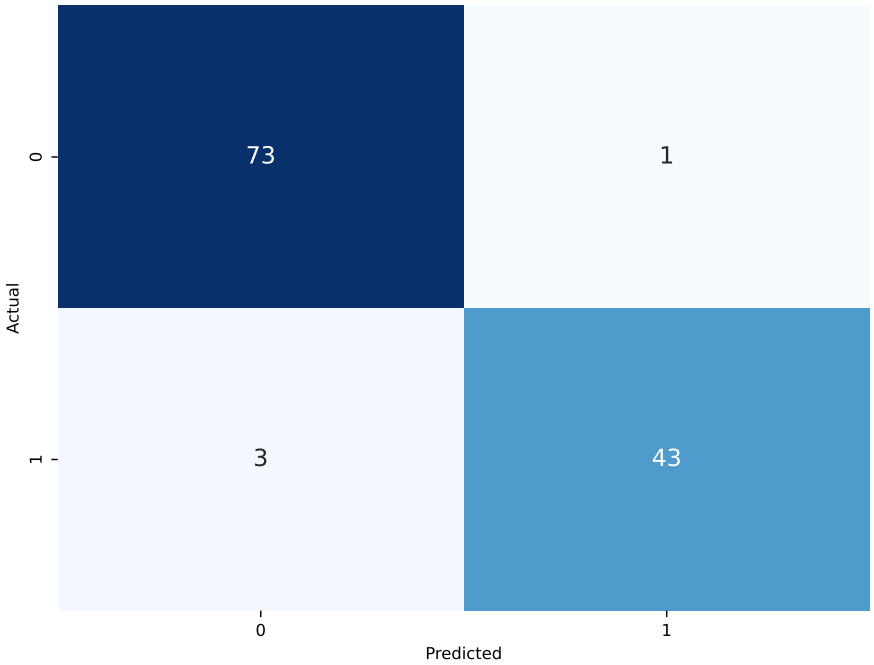

Fig. 3 Confusion Matrix for the Hybrid Quantum Neural Network

4.1 Performance Criteria

The performance metrics that determined the quality of the models' predictions of patients' outcomes from this study were: Precision, Recall, Accuracy, and F1-score.

Precision: Precision, or positive predictive value, here is the ratio of all patients actually with CKD to all the patients predicted with CKD (true positive and false positive)

Recall: It is also known as sensitivity, and it is the ratio of actual number of CKD patients that are correctly identified to the total number of patients with CKD.

F1-score: It is a metric that evaluates the accuracy of a classification model by considering both precision and recall. Specifically, it is the harmonic mean of precision and recall.

Accuracy: It refers to the proportion of correct predictions made by a model compared to the total number of predictions. It is calculated as the ratio of the number of correctly predicted instances (both true positives and true negatives).

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{1}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{2}$$

$$\text{F1-Score} = 2 \times \frac{\text{Recall} \times \text{Precision}}{\text{Recall} + \text{Precision}} \tag{3}$$

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \tag{4}$$

4.2 Models Benchmark

After running 10 evaluation, for the hybrid quantum-classical neural network, we noticed that the hybrid classifier performed really well across all evaluation metrics in terms of Accuracy, F1, and Recall, and made predictions as accurately as its classical counterparts trained on the same dataset [9], [22] while being run on a noisy simulated quantum processor.

Table 4 Benchmark of quantum models against classical counterparts

Model	Accuracy	F-Score	Precision	Recall
HQNN (Proposed)	97.50% (±0.72)	96.70% (±0.98)	97.78% (±0.04)	95.65% (±1.88)
Hossain et.al[10]	87.56%	89.58%	97.72%	82.69%
Parthasarathi et.al[12]	99.20%	98.90%	98.30%	98.50%
Qin et.al[27]	99.83%	99.86%	99.80%	99.84%
Khan et.al[15]	99.75%	99.80%	99.80%	99.80%

*Please note that the classical counterparts' performance metrics taken from [15], [27] and standalone quantum counterparts from [10], [12] are the highest obtained in each respective paper in terms of Accuracy.

4.3 Classification Report

We evaluated the performance of the hybrid quantum neural network in classifying the two distinct categories within the Chronic Kidney Disease (CKD) dataset—namely, 'ckd' (class 0) and 'notckd' (class 1). This assessment will be based on key performance metrics, including Precision, Recall, and F1-Score, computed separately for each class. These metrics provide a deeper understanding of how effectively the model identifies true positive cases for each category, balances false positives and false negatives, and maintains overall classification quality.

Table 5 Classification report for the hybrid quantum neural network

	F-Score	Precision	Recall	Support
Class 0 (ckd)	97.99%	97.34%	98.65%	74
Class 1 (notckd)	96.69%	97.78%	95.65%	46
Macro Avg.	97.34%	97.56%	97.15%	120
Weighted Avg.	97.45%	97.51%	97.50%	120
Accuracy			97.50%	

*Note: Accuracy is a standalone global score; it is not placed under the Recall column.

The hybrid quantum neural network demonstrated excellent performance, with a high overall accuracy of 97.50% and near-perfect precision, recall, and F1-scores for both classes. It correctly classified cases with minimal misclassification or false positives. Despite the slight imbalance of the dataset, the model remains well-balanced and performs consistently across both classes.

4.4 Probabilistic Risk Estimation & Feature Attribution Result

We sampled a batch of 16 random patient records from our test split, and ran it through our hybrid quantum neural network to give the probabilistic risk estimation and also passed it through the `IntegratedGradient` data attribution layer to get the biomarkers in this UCI CKD cross-sectional cohort that pushed a diagnosis towards the predicted class.

Table 6 Mean Feature Attribution of 16 Samples & Risk Estimation

Features	Total Mass
Albumin (al)	27.96%
Red Blood Cells (rbc)	25.79%
Hypertension (htn)	13.39%
Specific Gravity (sg)	7.05%
Diabetes Mellitus (dm)	4.79%
Packed Cell Volume (pcv)	3.95%
Blood Glucose Random (bgr)	3.38%
Hemoglogin (hemo)	3.31%
Serum Creatinine (sc)	2.98%
Blood Urea (bu)	2.76%
Pus Cell (pc)	2.47%
Red Blood Cell Count (rbcc)	2.18%

4.5 Model Interpretability & Circuit Expressivity

We further went ahead to evaluate the internal mechanics of our hybrid quantum neural network architecture. We looked at some key areas; the classical bottleneck layer for feature representation, circuit expressivity and effective dimensions, quantum layer bypass and layer feature ablation.

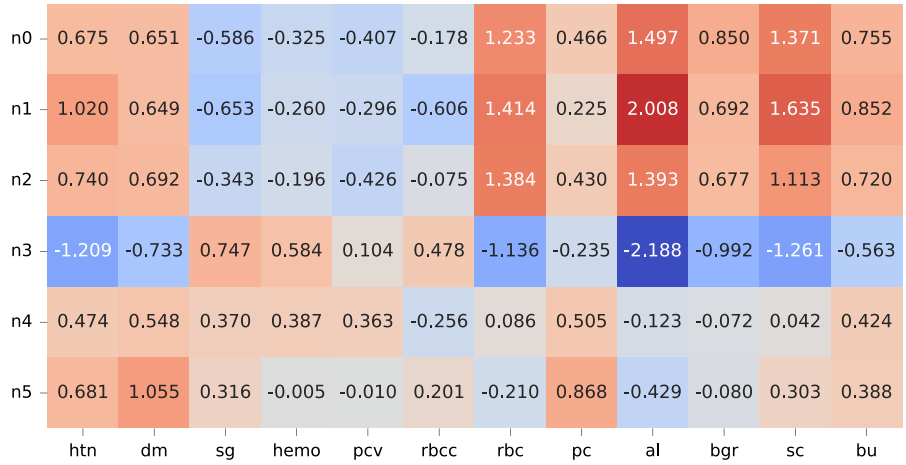

Fig. 4 Layer 1 Weights (Linear with Sigmoid Activation)

The first layer is a Linear layer with non-linear (Sigmoid) activation function, that compresses our 12 features into 6 features, and then into 3 dimensional latent space for quantum encoding. We see the primary group (neurons $n0$–$n2$) acting as aligned feature detectors, with dominant clinical drivers (al, rbc, sc, htn, dm, and bu) reinforcing primary disease indicators, while the inverted signal group (neuron $n3$) serves as antagonistic detectors assigning negative weights to risk markers (al, rbc, sc, htn, dm, bgr, bu, pc) while positively weighting protective features (sg, hemo, pcv, and rbcc) reflecting inverse clinical relationships with chronic kidney risk. This asymmetric relationship gives clear decision boundaries, preventing false-positive risk leakage that might wrongly predict an otherwise healthy patient as high risk.

We determined the quantum circuit expressivity by computing the Kullback-Leibler (KL) divergence between the circuits sampled fidelity distribution and the theoretical Haar fidelity distribution [34].

$$D_{\mathrm{KL}} = \int_0^1 P_{\mathrm{PQC}}(F)\, \ln\!\left(\frac{P_{\mathrm{PQC}}(F)}{P_{\mathrm{Haar}}(F)}\right) dF \tag{5}$$

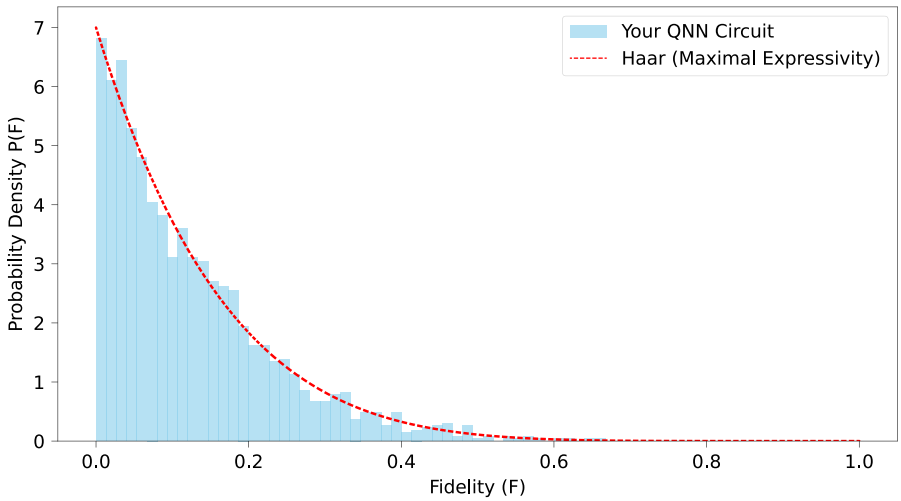

Fig. 5 Quantum Circuit Expressivity (KL Divergence)

We evaluated the functional capacity and resilience of our 3-qubit parameterized variational quantum circuit using Haar uniformity and effective dimension analysis. The circuit expressivity closely matches the theoretical Haar-random distribution, confirming high expressivity and uniform coverage of the Hilbert space in this particular setup. The circuit achieved a global effective dimension using **6.30 (6.27)** out of **9** trainable parameters, yielding a normalized capacity of **70.00% (69.72%)** on a sample size n where $n = \mathbf{1000}$; showing high parameter efficiency and expressive representation while remaining resilient against barren plateau during hybrid gradient backpropagation in this hybrid setup.

We bypassed the quantum layer completely to determine how much the feature attribution shifts, and we noticed that the model falls back on simple linear relationships to compensate for loss of non-linear representation.

Table 7 Quantum Layer Bypass Feature Attribution of 16 Samples

Features	Base	Bypass	ShiftΔ
Albumin (al)	27.96%	32.38%	+4.42%
Red Blood Cells (rbc)	25.79%	22.50%	−3.28%
Hypertension (htn)	13.39%	9.23%	−4.16%
Specific Gravity (sg)	7.05%	10.13%	+3.07%
Diabetes Mellitus (dm)	4.79%	1.00%	−3.79%
Packed Cell Volume (pcv)	3.95%	3.05%	+0.58%
Blood Glucose Random (bgr)	3.38%	3.99%	+0.62%
Hemoglogin (hemo)	3.31%	4.90%	+1.59%
Serum Creatinine (sc)	2.98%	3.03%	+0.05%
Blood Urea (bu)	2.76%	2.21%	−0.55%
Pus Cell (pc)	2.47%	3.05%	+0.58%
Red Blood Cell Count (rbcc)	2.18%	2.63%	+0.44%

The total absolute shift showed a **23.13%** attribution redistribution when we bypassed the quantum layer. To further solidify the impact of the quantum layer we ran a zeroing-out layer feature ablation. This yielded a mean layer feature ablation score of **0.2003** confirming that the quantum layer played an active non-trivial role in classification, accounting for approximately **20%** of the model's decision magnitude.

5 Conclusion & Future Works

In this study, the 12 best features were determined using an ensemble association test, features selected utilizing the Pearson Correlation Coefficient and Cramer's V Test were used to train and test the hybrid quantum-classical neural network. The hybrid quantum-classical neural network avoids costly computation altogether; it uses the MLP technique, which makes it faster for both training and testing in this case. In addition to this, the MLP technique had quantum layers with a variational quantum circuit that only used 3 qubits, whereas something like a QSVC needs n qubits ($n = 12$), thereby resulting in quadratic time complexity.

The main objective of this study was to leverage quantum algorithms to perform machine learning tasks using a hybrid quantum-classical neural network that can be run resiliently on current NISQ computers, which continues to see rapid use and improvements in fields like finance, pharmaceuticals, machine learning, and more. It is also meant to be a mode of entry in the real-world for future quantum computers [27], [28].

We were limited to running the hybrid quantum neural network on simulated quantum hardware. We were still able to get value from the quantum algorithms, giving us a path to resilient NISQ era applications. With the current research and development of fault-tolerant quantum computers, as they become readily available, their application in the field of healthcare – most notably in predictive medical diagnosis, will be highly valuable.

With the least amount of features (12 out of 24), the model was able to make accurate predictions. Features selected through association tests; hemoglobin (hemo), specific gravity (sg), hypertension (htn), diabetes mellitus (dm), packed cell volume (pcv), red blood count (rbc), albuminuria (al) and then the domain-guided relevant features like serum creatinine (sc) and blood urea (bu). This maintained medically relevant features in the final set, despite internal collinearities which were resolved using our classical layer to compress them into low dimensional latent space.

We believe that given that the original dataset contained only 400 samples – the volume of data was relatively small. As a result, we do not consider it sufficient to reliably assess the effectiveness or generalizability of the model; while the proposed model performed really well on this set of cross-sectional patient data, its sensitivity remains unverified at scales typically associated with big data.

Finally, with the rapid adoption of quantum machine learning algorithms in other fields like high-energy physics [38] and their application in the field of healthcare, studies like this can be taken further by adopting solutions seen in [3] using the techniques in [36] or furthermore developing ensemble techniques like [37] but applied to quantum machine learning algorithms. The clear path we see to more adoption

and a strong feedback loop in quantum machine learning especially in the field of healthcare is designing hybrid architectures and putting them to work in the real-world with robust large-scale datasets.

6 Declarations

The authors make the following declarations regarding funding, data, authorship, and other relevant aspects of this work.

6.1 Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

6.2 Conflicts Of Interest / Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

6.3 Availability Of Data and Materials

The dataset used in this study is publicly and readily available and can be accessed through the UCI Machine Learning Repository under the title "Chronic Kidney Disease Dataset" with ID `336` (Rubini, L., Soundarapandian, P., & Eswaran, P., 2015). (https://doi.org/10.24432/C5G020)

6.4 Authors' Contributions

Emmanuel Ismaila Watila conceptualized the study, designed the hybrid quantum neural network architecture, performed the experiments, and drafted the manuscript. Peter Yusuf Dibal & Stephen Joseph Bassi contributed to the literature review, validation, and manuscript editing, and assisted in recalling and clarifying methodological and analytical details during the research and writing process. All authors read and approved the final manuscript.

6.5 Acknowledgement

We would also like to appreciate Rubini, L., Soundarapandian, P., Eswaran, P., & University of California, Irvine, for making the CKD dataset openly accessible. Providing this valuable resource has played a crucial role in enabling researchers, students, engineers, enthusiasts, and more — to explore real-world medical data to develop innovative solutions aimed at improving healthcare outcomes in the future.

References

[1] K. J. Jager and S. D. Fraser. The ascending rank of chronic kidney disease in the global burden of disease study. *Nephrology Dialysis Transplantation*, 32(suppl. 2):ii121–ii128, 2017.

[2] T. O. Olanrewaju, A. Aderibigbe, A. A. Popoola, K. T. Braimoh, M. O. Buhari, O. T. Adedoyin, and K. Klipstein-Grobusch. Prevalence of chronic kidney disease and risk factors in North-Central Nigeria: a population-based survey. *BMC Nephrology*, nov 2020.

[3] A. M. Athavale, P. D. Hart, M. Itteera, D. Cimbaluk, T. Patel, A. Alabkaa, and H. Kulkarni. Development and Validation of a Deep Learning Model to Quantify Interstitial Fibrosis and Tubular Atrophy From Kidney Ultrasonography Images. *JAMA Network Open*, 4(5):e2111176, 2021.

[4] M. Wojciechowska, S. Malacrino, N. Garcia Martin, H. Fehri, and J. Rittscher. Early Detection of Liver Fibrosis Using Graph Convolutional Networks. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, pages 217–226, 2021.

[5] H. Wang, T. Fu, Y. Du, W. Gao, K. Huang, Z. Liu, and M. Zitnik. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, 2023.

[6] F. Zhu, P. Xu, and J. Zong. Moore's Law: The potential, limits, and breakthroughs. *Applied and Computational Engineering*, 10(1):307–315, 2023.

[7] Q. A. Memon, M. A. Ahmad, and M. Pecht. Quantum Computing: Navigating the future of computation, challenges, and technological breakthroughs. *Quantum Reports*, 6(4):627–663, 2024.

[8] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer. Elucidating reaction mechanisms on quantum computers. *Proceedings of the National Academy of Sciences*, 114(29):7555–7560, 2017.

[9] J. D. Martin-Guerrero and L. Lamata. Quantum Machine Learning: A tutorial. *Neurocomputing*, 470:457–461, 2021.

[10] M. M. Hossain, A. Munir, M. A. Habib, M. S. Iqbal, M. M. Hasan, and M. M. Rahman. Performance analysis of classical and quantum support vector machines for diagnosis of chronic kidney disease. *Informatics and Health*, 2(2):179–193, 2025.

[11] D. Arthur and P. Date. A Hybrid Quantum-Classical Neural Network Architecture for Binary Classification, jan 2022.

[12] P. Parthasarathi, H. M. Alshahrani, K. Venkatachalam, and J. Cho. Variational quantum classifier-based early identification and classification of chronic kidney disease using sparse autoencoder and LASSO shrinkage. *PeerJ Computer Science*, 11:e2789, 2025.

[13] P. J. Kiruthiga and H. Stanly. Quantum-Enhanced SVM for Chronic kidney Disease prediction. In *Machine Intelligence for Research and Innovations*, volume

1515 of *Lecture Notes in Networks and Systems*, pages 297–308. Springer, 2026.

[14] N. Deepika, P. RaniAnnapurna, T. S. Kavyanjali, V. Mahesh, and M. Sriramulu. Early detection of chronic diseases using quantum machine learning algorithms. *SPU Journal of Science Technology and Management Research*, pages 161–167, 2026.

[15] R. S, B. Bharathi, P. Jeyanthi, and M. Ramesh. Chronic Kidney Disease Prediction using Machine Learning Models. *International Journal of Engineering and Advanced Technology*, 9(1):6364–6367, 2019.

[16] B. Khan, R. Naseem, F. Muhammad, G. Abbas, and S. Kim. An Empirical evaluation of machine learning techniques for chronic kidney disease prophecy. *IEEE Access*, 8:55012–55022, 2020.

[17] M. M. Nishat, F. Faisal, R. R. Dip, S. M. Nasrullah, R. Ahsan, F. Shikder, and M. A. Hoque. A comprehensive analysis on detecting chronic kidney disease by employing machine learning algorithms. *EAI Endorsed Transactions on Pervasive Health and Technology*, 7(29):e1, 2021.

[18] C. Sabanayagam, D. Xu, D. S. W. Ting, S. Nusinovici, R. Banu, H. Hamzah, and T. Y. Wong. A deep learning algorithm to detect chronic kidney disease from retinal photographs in community-based populations. *The Lancet Digital Health*, 2(6):e295–e302, 2020.

[19] M. Almasoud and T. E. Detection of Chronic Kidney Disease using Machine Learning Algorithms with Least Number of Predictors. *International Journal of Advanced Computer Science and Applications*, 10(8), 2019.

[20] L. Rubini, P. Soundarapandian, and P. Eswaran. Chronic Kidney Disease [Dataset]. UCI Machine Learning Repository, dataset, 2015.

[21] Danilo Schmidt, Matthias Niemann, and Gabriela Lindemann von Trzebiatowski. The Handling of Missing Values in Medical Domains with Respect to Pattern Mining Algorithms. In *Proceedings of the International Workshop on Concurrency, Specification and Programming (CS&P 2015)*, volume 1492 of *CEUR Workshop Proceedings*, pages 147–154, 2015.

[22] M. J. Azur, E. A. Stuart, C. Frangakis, and P. J. Leaf. Multiple imputation by chained equations: what is it and how does it work? *International Journal of Methods in Psychiatric Research*, 20(1):40–49, 2011.

[23] A. H. Ibrahim, O. C. Karabulut, B. A. Karpuzcu, E. Türk, and B. E. Süzek. A correlation coefficient-based feature selection approach for virus-host protein-protein interaction prediction. *PLoS ONE*, 18(5):e0285168, 2023.

[24] S. Lopez-Giacoman. Biomarkers in chronic kidney disease, from kidney function to kidney damage. *World Journal of Nephrology*, 4(1):57, 2015.

[25] A. Gupta, T. Sontakke, S. Acharya, and S. Kumar. A Comprehensive review of biomarkers for chronic kidney disease in Older Individuals: Current perspectives and future directions. *Cureus*, 2024.

[26] Goce Dimeski and Oliver Treacy. Biochemical Tests for Diagnosing and Evaluation Stages of Chronic Kidney Disease. In Ane Claudia Fernandes Nunes, editor, *Chronic Kidney Disease – Beyond the Basics*. IntechOpen, jul 2022.

[27] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf. Training deep quantum neural networks. *Nature Communications*, 11(1), 2020.

[28] J. Qin, L. Chen, Y. Liu, C. Liu, C. Feng, and B. Chen. A Machine learning Methodology for diagnosing Chronic kidney Disease. *IEEE Access*, 8:20991–21002, 2019.

[29] Y. Amirgaliyev, S. Shamiluulu, and A. Serek. Analysis of Chronic Kidney Disease Dataset by Applying Machine Learning Methods. In *2018 IEEE 12th International Conference on Application of Information and Communication Technologies (AICT)*, 2018.

[30] F. F. Kaze, M. P. Kowo, I. N. Wagou, M. Maimouna, H. D. M. E. Fouda, and M. P. Halle. Hematological Disorders during Chronic Kidney Disease Stages 3 to 5 Non-Dialysed in Cameroon. *Open Journal of Nephrology*, 10(2):61–72, 2020.

[31] S. Uchino, R. Bellomo, and D. Goldsmith. The meaning of the blood urea nitrogen/creatinine ratio in acute kidney injury. *Clinical Kidney Journal*, 5(2):187–191, 2012.

[32] Y. Cao, J. Romero, and A. Aspuru-Guzik. Potential of quantum computing for drug discovery. *IBM Journal of Research and Development*, 62(6):6:1–6:20, 2018.

[33] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. F. Bulmer, and J. Lavoie. Quantum computational advantage with a programmable photonic processor. *Nature*, 606(7912):75–81, 2022.

[34] V. Bonnici. Kullback-Leibler divergence between quantum distributions, and its upper-bound, aug 2008.

[35] V. Jackins, S. Vimal, M. Kaliappan, and M. Y. Lee. AI-based smart prediction of clinical disease using random forest classifier and Naive Bayes. *The Journal of Supercomputing*, 77(5):5198–5219, 2020.

[36] J. Landman, N. Mathur, Y. Y. Li, M. Strahm, S. Kazdaghli, A. Prakash, and I. Kerenidis. Quantum Methods for Neural Networks and Application to Medical Image Classification. *Quantum*, 6:881, 2022.

[37] L. A. Ha and E. Mohamed. Combining Text and Images for Film Age Appropriateness Classification. *Procedia Computer Science*, 189:242–249, 2021.

[38] H. Yin, Z. Hu, H. Zeng, J. Guan, and J. Wang. Application of quantum machine learning using variational quantum classifier in accelerator physics, jun 2025.