

Quantum enhanced machine learning for medical image analysis: a hybrid approach

Applied
Computing and
Informatics

Soha Rawas

*Department of Mathematics and Computer Science, Faculty of Science,
Beirut Arab University, Beirut, Lebanon, and*

Duaa AlSaeed

*College of Computer and Information Sciences, King Saud University,
Riyadh, Saudi Arabia*

Received 10 October 2025
Revised 31 October 2025
12 December 2025
Accepted 25 December 2025

Abstract

Purpose – This study investigates the integration of quantum-enhanced machine learning (QML) into medical image analysis, focusing on classification and feature extraction tasks. It aims to evaluate whether hybrid quantum-classical models can overcome the limitations of traditional machine learning in handling high-dimensional, clinically relevant imaging data.

Design/methodology/approach – Hybrid architectures combining classical preprocessing with quantum components, specifically variational quantum circuits (VQCs) and quantum kernel methods, were developed and implemented in simulated quantum environments using Qiskit and PennyLane on Google Colab. Publicly available datasets (ChestX-ray14, Brain Tumor MRI and CT Hemorrhage) were preprocessed and encoded into quantum states. Model performance was evaluated against classical baselines (MLP and CNN) using accuracy, precision, recall, F1-score, convergence speed and trainable parameters.

Findings – Quantum-enhanced models achieved competitive or superior performance compared to classical models with significantly fewer trainable parameters and faster convergence. The Quantum Kernel SVM demonstrated the highest precision (89.1%) and F1-score (88.5%) using only 8 parameters, while VQC models converged up to $2\times$ faster than classical baselines. These results highlight the potential of QML to improve efficiency, robustness and scalability in diagnostic imaging workflows, particularly in resource-constrained or data-limited settings.

Research limitations/implications – The study relies on simulated quantum environments, which do not capture hardware noise or qubit limitations. Future work will extend to real quantum devices, multimodal data integration and interpretability strategies.

Practical implications – Quantum kernel methods and lightweight VQC architectures enable faster, resource-efficient diagnostic modeling, making them suitable for clinical environments with limited computational capacity or time-sensitive decision-making needs.

Social implications – Improving the efficiency and accessibility of medical image analysis has the potential to support faster diagnosis and enhance healthcare delivery, particularly in underserved or resource-limited regions.

Originality/value – This study provides a practical evaluation of hybrid quantum-classical models applied to multiple real medical imaging datasets in a reproducible simulated environment. It offers an early-stage framework demonstrating how QML techniques can be integrated into existing machine learning pipelines for medical image analysis.

Keywords Quantum enhanced machine learning, Medical imaging diagnostics, Hybrid quantum-classical algorithms, Quantum data analysis, Artificial intelligence in healthcare

Paper type Research article

© Soha Rawas and Duaa AlSaeed. Published in *Applied Computing and Informatics*. Published by Emerald Publishing Limited. This article is published under the Creative Commons Attribution (CC BY 4.0) licence. Anyone may reproduce, distribute, translate and create derivative works of this article (for both commercial and non-commercial purposes), subject to full attribution to the original publication and authors. The full terms of this licence may be seen at [Link to the terms of the CC BY 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Funding declaration: The work was supported by a grant from the Research center of College of Computer and Information Sciences”, Deanship of Scientific Research, King Saud University, Riyadh, Saudi Arabia.

Competing interest: The authors declare that they have no competing interests.



Applied Computing and Informatics
Emerald Publishing Limited
e-ISSN: 2210-8327
p-ISSN: 2634-1964
DOI 10.1108/ACI-10-2025-0448

Nomenclature

AI	Artificial Intelligence
CNN	Convolutional Neural Network
CT	Computed Tomography
MLP	Multilayer Perceptron
MRI	Magnetic Resonance Imaging
QML	Quantum-enhanced Machine Learning
SVM	Support Vector Machine
VQC	Variational Quantum Circuit

1. Introduction

Medical imaging plays a central role in modern clinical practice, enabling early diagnosis, treatment planning, and disease monitoring across a wide range of conditions [1]. As imaging modalities such as magnetic resonance imaging (MRI), computed tomography (CT), and digital pathology continue to evolve, they produce increasingly complex, high dimensional datasets [2]. These datasets demand robust computational approaches capable of extracting meaningful insights with accuracy, efficiency, and reliability.

In recent years, artificial intelligence (AI), particularly deep learning, has made significant progress in automating medical image interpretation [3, 4]. However, conventional machine learning algorithms often face challenges related to scalability, data imbalance, overfitting, and computational demands, especially when working with limited labeled data or large scale imaging repositories [5].

Quantum computing offers a promising complementary framework, introducing fundamentally different approaches to computation based on principles such as superposition and entanglement [6]. Superposition allows quantum bits (qubits) to exist in multiple states simultaneously, while entanglement creates strong correlations between qubits that can help with complex pattern recognition. Within this context, quantum enhanced machine learning (QML) has emerged as a hybrid approach that combines classical deep learning models with quantum components such as variational quantum circuits and quantum kernels [7]. These hybrid models have demonstrated potential in addressing complex optimization tasks, accelerating convergence, and improving generalization in high dimensional data scenarios by exploiting the exponential computational space of quantum systems.

Recent studies have increasingly explored the integration of QML into biomedical and imaging applications. For instance, hybrid quantum classical models have been employed for enhancing diagnostic precision and accelerating feature extraction from MRI and CT datasets [8, 9]. These works demonstrate the potential of QML to improve classification performance and computational efficiency in healthcare. However, most existing studies focus on theoretical or small scale implementations, leaving a gap in demonstrating reproducible, scalable QML frameworks for real world medical imaging pipelines. The present work addresses this gap by proposing a hybrid, simulation based approach that evaluates both variational and kernel based QML models under consistent experimental conditions.

This paper investigates the application of QML techniques to medical image analysis, with a particular focus on diagnostic tasks such as classification and feature extraction. The hybrid approach combines classical preprocessing for reducing dimensionality with quantum circuits to handle the computationally demanding tasks of pattern recognition. Rather than emphasizing visual reconstruction, the aim is to explore how quantum enhanced models can support efficient, scalable, and analytically robust decision making processes. All implementations are carried out in simulated environments using Qiskit and PennyLane within Google Colab, ensuring reproducibility and practical accessibility without the need for physical quantum hardware.

The selection of this research focus is motivated by two key factors. First, classification and feature extraction are fundamental pillars of diagnostic radiology and pathology, where improved accuracy and efficiency can have direct clinical impact. Second, these tasks, which often involve mapping high dimensional data to decision boundaries, are theoretically well suited to the purported advantages of quantum computing, such as operating in high dimensional Hilbert spaces and efficiently handling certain types of complex optimization. Therefore, investigating QML for these specific tasks provides a critical and relevant testbed for assessing its practical utility in healthcare.

This study is guided by the following research question: “Can hybrid quantum classical machine learning models provide a scalable, efficient, and accurate framework for medical image classification and feature extraction?”

The primary objectives of this work are:

- (1) To design a hybrid quantum classical pipeline that integrates classical preprocessing with quantum enhanced models, specifically Variational Quantum Circuits (VQCs) and Quantum Kernel Methods.
- (2) To rigorously evaluate the performance of these models against classical baselines (CNNs and MLPs) on diverse, publicly available medical imaging datasets.
- (3) To analyze the potential advantages of QML, such as convergence speed, model complexity, and robustness, within a reproducible, simulation based environment.

The remainder of this paper is organized as follows: [Section 2](#) presents the background and related work in both classical AI and quantum machine learning for healthcare. [Section 3](#) details the methodology, outlining the proposed hybrid framework, quantum components, and experimental pipeline. [Section 4](#) discusses the results and insights from the evaluation of quantum enhanced models. Finally, [Section 5](#) concludes the study and outlines future research directions.

2. Background and related work

The intersection of artificial intelligence and medical imaging has seen major growth over the past decade [10, 11]. Deep learning models, especially convolutional neural networks (CNNs), have performed well in tasks like tumor detection, organ segmentation, and disease classification, with some systems achieving radiologist level performance in specific domains. However, these models often need large amounts of labeled data, high computational power, and careful tuning [12]. They also struggle to generalize across different imaging types and patient groups.

To overcome these limitations, research has expanded into alternative machine learning paradigms, including self-supervised learning, federated learning, and generative modeling [13–15]. While these methods address specific challenges, such as privacy preservation or data augmentation, the scalability and interpretability of models remain critical concerns, particularly in real time or resource constrained clinical environments.

QML introduces a new computational perspective, leveraging quantum mechanical phenomena to encode and process information [16]. These phenomena include superposition (where quantum systems can exist in multiple states at once), entanglement (non-local connections between quantum particles), and interference (the way quantum states combine to increase or decrease certain probabilities). Notably, hybrid quantum classical models [17], such as those combining classical neural networks with parameterized quantum circuits, are designed to operate on near term quantum devices or their simulated equivalents. Variational quantum circuits (VQCs) and quantum kernels are among the most promising techniques, allowing for expressive and compact model representations that may outperform classical counterparts on certain structured data problems [18].

Several preliminary studies have demonstrated the feasibility of applying QML to biomedical domains, including genomics, drug discovery, and radiomics [19]. In the context of medical imaging, quantum enhanced methods have shown potential for improving feature space separation and reducing training times [20]. However, much of the existing work remains conceptual or limited to small-scale datasets due to hardware constraints and a lack of standardized workflows.

This paper builds on these foundational efforts by focusing on the application of QML for practical and scalable image analysis tasks, using simulated quantum environments that replicate the behavior of quantum hardware. The aim is to explore the analytical advantages of quantum-enhanced models and evaluate their utility in supporting more efficient and accurate clinical decision-making.

3. Methodology: hybrid framework for medical image classification and feature extraction

This study investigates the application of hybrid quantum-classical machine learning models to medical image classification and feature extraction tasks. The framework is not intended for segmentation, registration, or image enhancement, but rather focuses on diagnostic decision-making based on extracted imaging features. The methodology integrates quantum-enhanced components within conventional machine learning pipelines to evaluate their analytical advantages in handling high dimensional and complex clinical data.

3.1 Overview of the proposed framework

The proposed framework adopts hybrid architecture that combines classical preprocessing and feature engineering stages with quantum enhanced model training. Initially, medical images undergo standard preprocessing techniques, including resizing, normalization, and grayscale conversion, to reduce dimensionality while preserving diagnostically significant features. For the ChestX-ray14 dataset, images were resized to 224×224 pixels, normalized to $[0,1]$ range, and converted to grayscale. MRI and CT images underwent additional contrast enhancement using histogram equalization. These processed feature vectors are subsequently encoded into quantum states using suitable encoding strategies and fed into quantum machine learning models for classification. Figure 1 illustrates the overall structure of the hybrid system and the interaction between classical and quantum components.

In this framework, the evaluation process begins with dataset preprocessing and feature encoding, followed by the training of two quantum models, Variational Quantum Circuit (VQC) and Quantum Kernel SVM, each implemented through PennyLane and Qiskit simulators, respectively. The performance of these models is then assessed using metrics such as accuracy, precision, recall, and F1-score, and compared directly with classical baselines (CNN and MLP). This process, represented in Figure 1, captures the complete workflow from input acquisition to model benchmarking.

This workflow enables the integration of quantum components at different levels of abstraction, allowing for systematic comparison with classical counterparts in terms of performance, convergence behavior, and generalization capabilities.

3.2 Quantum formalism

The hybrid framework leverages two primary quantum approaches, formally defined below.

Variational Quantum Circuit (VQC) Model: The VQC operates as a parameterized quantum model for classification. The predicted output for a given data input $\mathbf{x}$ is given by the expectation value:

$$f(\mathbf{x}; \boldsymbol{\theta}) = \langle \mathbf{0} | U^\dagger(\mathbf{x}, \boldsymbol{\theta}) M U(\mathbf{x}, \boldsymbol{\theta}) | \mathbf{0} \rangle$$




Figure 1. Block diagram of the hybrid quantum–classical framework, illustrating data preprocessing, quantum encoding, model training, and evaluation stages within the proposed QML workflow

where:

- (1) $U(\mathbf{x}, \boldsymbol{\theta})$ is the parameterized quantum circuit, composed of a feature embedding layer $S(\mathbf{x})$ that encodes the classical data, followed by a variational layer $W(\boldsymbol{\theta})$ with tunable parameters $\boldsymbol{\theta}$.
- (2) M is a measurement operator, typically the Pauli-Z operator on one or more qubits.
- (3) The parameters $\boldsymbol{\theta}$ are optimized iteratively using a classical optimizer to minimize a loss function (e.g. mean squared error or cross-entropy).

Quantum Kernel Method: The quantum kernel leverages a quantum feature map $\boldsymbol{\varphi}(\mathbf{x})$ to project classical data into a high-dimensional Hilbert space. The quantum kernel matrix elements between two data points $\mathbf{x}_i$ and $\mathbf{x}_j$ are defined as:

$$K_{ij} = |\langle \boldsymbol{\varphi}(\mathbf{x}_j) | \boldsymbol{\varphi}(\mathbf{x}_i) \rangle|^2$$

This kernel, K_{ij} , quantifies the overlap between the two quantum states and is used within a classical Support Vector Machine (SVM) to construct the decision boundary. The kernel is computed on a quantum processor (or simulator) by evaluating the transition probability between the two feature states.

3.3 Detailed implementation pipeline

To ensure reproducibility, we detail the sequential steps that transform a raw medical image into a format processable by our quantum models, with specific parameter choices.

- (1) **Image Preprocessing and Feature Reduction:** All images are resized to 224×224 pixels and converted to grayscale, resulting in a 50,176-dimensional vector per image. To make this compatible with near-term quantum simulations, we apply Principal Component Analysis (PCA) to the flattened image vectors. We retain the top 256

principal components, which capture 95% of the cumulative variance in the training set across all datasets. These 256-dimensional vectors are then min-max normalized to the range $[0, 1]$.

- (2) Classical to Quantum Encoding: The reduced classical vector $x \in \mathbb{R}^{256}$ is encoded into the state of $n = 8$ qubits using two distinct strategies:
 - Angle Encoding (for VQC): The first 8 features of x are used. For qubit i , the feature x_i is encoded as a rotation angle via the gate $R_Y(x_i \cdot \pi)$.
 - Amplitude Encoding (for Quantum Kernel): The full normalized 256-dimensional vector is encoded into the $2^8 = 256$ computational basis amplitudes of an 8-qubit state. This is implemented using the *AmplitudeEmbedding* template in PennyLane.
- (3) VQC Architecture: Our VQC employs the following explicit structure:
 - Encoding Layer: An angle encoding layer as described above.
 - Variational Ansatz: A hardware-efficient Ansatz with 3 repeated layers. Each layer consists of:
 - Single-qubit $R_Y(\theta)$ rotations on all 8 qubits.
 - Entangling CNOT gates applied in a circular pattern (qubit i to qubit $(i + 1) \bmod 8$).
 - Measurement: The expectation value of the Pauli-Z operator on the first qubit is measured.
The total number of trainable parameters is 18: 8 parameters from the initial encoding layer, plus 3 layers $\times$ 3 parameterized gates per qubit in a specific configuration that yields 10 additional parameters. The circuit schematic is shown in [Figure 2](#).

In summary, the VQC architecture is defined by: 8 qubits, a depth of 3 variational layers, circular entanglement, and a gate set of parameterized R_Y rotations and CNOT gates.

- (4) Quantum Kernel Estimation: For the Quantum Kernel SVM, the kernel matrix K is computed using the 8-qubit amplitude-encoded states. Each element $K_{ij} = |\langle \phi(x_j) | \phi(x_i) \rangle|^2$ is estimated using the statevector simulator in Qiskit. The resulting kernel matrix is used with a classical linear SVM from scikit-learn for classification.




Figure 2. Schematic of the 8-qubit VQC architecture. The circuit comprises an encoding layer (EN) implementing angle encoding via $R_Y(x_i \cdot \pi)$ gates, followed by three variational layers (L1, L2, L3), each containing parameterized $R_Y(\theta)$ rotations and circular entanglement via CNOT gates (●—X connections). Measurement is performed on the first qubit (q0).

3.3.1 Classical-to-quantum encoding. After PCA reduction, each image is represented as a normalized 256-dimensional feature vector $x \in [0, 1]^{256}$. This classical vector is encoded into the state of $n = 8$ qubits using two distinct strategies, each suited to a specific quantum model.

Angle Encoding (for VQC): The first 8 features $\{x_1, \dots, x_8\}$ are used. For qubit i , the feature x_i is encoded as a rotation angle via an R_Y gate:

$$|\psi_i\rangle = R_Y(x_i \cdot \pi) |0\rangle = \cos\left(\frac{x_i \pi}{2}\right) |0\rangle + \sin\left(\frac{x_i \pi}{2}\right) |1\rangle.$$

The full 8-qubit encoded state is:

$$|\psi_{angle}\rangle = \bigotimes_{i=1}^8 R_Y(x_i \pi) |0\rangle^{\otimes 8}.$$

This encoding uses only 8 of the 256 features directly; the remaining features are implicitly represented via classical preprocessing and the variational layers that follow.

Amplitude Encoding (for Quantum Kernel): The full 256-dimensional vector x is encoded into the amplitudes of an 8-qubit state. Since $2^8 = 256$, the normalized vector x is placed directly into the computational basis amplitudes:

$$|\psi_{amp}\rangle = \sum_{k=1}^{256} x_k |k\rangle,$$

where $|k\rangle$ denotes the k -th computational basis state of 8 qubits. This is implemented using the *AmplitudeEmbedding* template in PennyLane, which requires the input vector to be normalized ($\|x\|_2 = 1$).

Qubit Count Justification: We use 8 qubits due to simulation constraints, which allow $2^8 = 256$ amplitude-encoded features. This matches the PCA-reduced dimension and maintains feasibility in current quantum simulators. The encoding circuits for both strategies are integrated into the model architectures as shown in [Figure 2](#).

3.3.2 VQC architecture and Ansatz specification. The VQC employs a hardware-efficient Ansatz with circular entanglement, designed for near-term quantum devices. The circuit consists of:

- (1) Encoding Layer: Angle encoding via $R_Y(x_i \pi)$ gates on each of the 8 qubits.
- (2) Variational Layers: Three repeated layers, each containing:
 - Parameterized $R_Y(\theta)$ rotations on all qubits.
 - CNOT gates applied in a circular topology (qubit i to qubit $(i + 1) \bmod 8$).
- (3) Measurement: Expectation value of the Pauli-Z operator on the first qubit.

Circuit Depth and Gate Count:

- (1) Total quantum depth: 4 layers (1 encoding + 3 variational).
- (2) Total gates: 32 R_Y rotations (8 encoding + 24 variational) and 24 CNOT gates.
- (3) Trainable parameters: 18 (8 from encoding angles, 10 from variational R_Y rotations).

The full circuit schematic is illustrated in [Figure 2](#).

3.4 Quantum components

Two distinct quantum enhanced modeling approaches are explored in this study:

- (1) *Variational Quantum Circuits (VQCs)*: Parameterized quantum circuits serve as trainable models capable of capturing complex data distributions. A classical optimization algorithm iteratively adjusts the circuit parameters to minimize a defined loss function, enabling supervised learning for multi-class image classification tasks.
- (2) *Quantum Kernel Methods*: Classical feature vectors are embedded into a high dimensional quantum Hilbert space through a quantum feature map. A quantum kernel matrix is then computed and employed within a classical support vector machine (SVM) or similar classifier to enhance class separability.

Both methods aim to exploit the representational power of quantum computing to achieve improved performance with reduced model complexity, particularly beneficial when working with limited or noisy medical imaging datasets.

3.5 Dataset description

To ensure clinical relevance and methodological rigor, experiments are conducted using publicly available medical imaging datasets, including:

- (1) *ChestX-ray14* [21]: A large scale dataset comprising frontal chest X-ray images labeled with 14 common thoracic pathologies.
- (2) *Brain Tumor MRI Dataset* [22]: A collection of MRI scans containing cases with and without brain tumors, used for binary classification tasks.
- (3) *NIH ChestX-ray14 Dataset* [21]: Provides multi label annotations for a variety of lung diseases.
- (4) *Kaggle CT-Scan Brain Hemorrhage Dataset* [23]: Contains axial CT slices annotated for the presence and type of intracranial hemorrhage.

These datasets were selected based on diagnostic complexity, modality diversity (X-ray, MRI, CT), and availability of standardized labels, ensuring broad applicability across different imaging modalities and clinical conditions. For training and evaluation, each dataset was balanced through controlled sampling to mitigate class imbalance, particularly for minority classes such as subtle hemorrhages and small tumors. Balancing was performed via random undersampling of the majority class within the training split only, ensuring the validation and test sets remained untouched and representative of the original class distribution. This protocol prevents data leakage and provides a realistic evaluation of model performance on imbalanced clinical data. Images were resized to 224×224 pixels, normalized, and converted to grayscale before encoding. For multi class datasets like ChestX-ray14, only the most diagnostically significant pathologies were selected to ensure robust classification. Each dataset was divided into 70% training, 15% validation, and 15% testing subsets, maintaining consistent preprocessing protocols across modalities.

3.6 Simulation environment

All quantum operations are simulated using:

- (1) *Qiskit* (IBM) [24]: Utilized for implementing quantum kernel matrices and interfacing with quantum SVMs.
- (2) *PennyLane* (Xanadu) [25]: Employed for constructing and optimizing variational quantum circuits, with seamless integration into classical deep learning libraries such as PyTorch and TensorFlow.

Experiments are executed in a cloud-based environment using Google Colab [26], which provides scalable computational resources and ensures reproducibility across multiple runs. This simulation approach introduces certain limitations compared to real quantum hardware,

including idealized noise-free operations and limited qubit counts, has some limits compared to real quantum hardware, like idealized, noise free conditions and fewer qubits, but it still provides a useful way to develop and test algorithms. This setup eliminates dependency on physical quantum hardware while closely simulating its behavior for early stage research and validation.

All quantum simulations utilized the statevector simulator backend, which computes the exact evolution of the quantum state. This choice is standard for algorithm development and validation in the NISQ era, as it provides noise-free results that clearly reveal the theoretical performance of the quantum circuits without the confounding factor of simulated sampling noise or hardware error, which are separate research challenges.

3.7 Experimental pipeline

The experimental workflow consists of the following structured stages:

- (1) **Dataset Acquisition and Preprocessing:** Medical images are resized, normalized, converted to grayscale, and partitioned into training, validation, and test subsets.
- (2) **Feature Encoding:** Processed feature vectors are encoded into quantum states using either angle encoding or amplitude encoding schemes, depending on the quantum model architecture.
- (3) **Model Training:** Hybrid quantum classical models are trained using simulated quantum backends. For VQC models, classical optimizers update circuit parameters to minimize cross-entropy loss; for kernel based models, quantum kernels are computed and used in conjunction with classical classifiers.
- (4) **Performance Evaluation:** Models are evaluated using key metrics such as accuracy, precision, recall, F1-score, training time, and convergence rate. Results are benchmarked against classical baselines, including convolutional neural networks (CNNs) and multilayer perceptrons (MLPs). The selection of simpler CNN and MLP architectures as classical baselines was intentional, designed to facilitate a parameter-matched and complexity-controlled comparison. Given the limited qubit count (8 qubits) and circuit depth feasible in current quantum simulations, we intentionally constrained the classical models to a comparable scale in terms of trainable parameters and architectural depth. This ensures that observed differences in performance, convergence, and efficiency are attributable to the computational paradigm (quantum vs. classical) rather than disparities in model capacity. While state-of-the-art classical models (e.g. ResNet, Vision Transformers) achieve higher absolute accuracy on well-established benchmarks, they require orders of magnitude more parameters and computational resources. Our aim is not to outperform such models in raw accuracy, but to demonstrate that quantum-enhanced models can achieve competitive diagnostic performance with significantly fewer parameters and faster convergence, a compelling advantage for resource-constrained clinical environments
- (5) **Analysis and Comparison:** Comparative analysis focuses on assessing improvements in convergence efficiency, robustness to data scarcity, and generalization across diverse imaging modalities.

4. Results and discussion

This section presents a detailed evaluation of the hybrid quantum classical models applied to medical image classification tasks. Experiments were conducted using simulated quantum

ACI environments on publicly available datasets, with emphasis on diagnostic accuracy, convergence efficiency, model complexity, and generalization capabilities.

All models were trained and evaluated under identical conditions for fairness: same dataset splits, preprocessing steps, and training epochs (up to 50). Quantum simulations were performed using Qiskit for kernel-based methods and PennyLane for VQCs, both integrated into Google Colab for reproducibility.

4.1 Performance metrics

The models were evaluated using the following standard classification metrics:

- (1) Accuracy: Ratio of correctly predicted samples over total samples.
- (2) Precision, Recall, F1-Score: To assess performance on imbalanced class distributions, especially relevant in medical imaging where pathological cases are often rare.
- (3) Training Time and Convergence Speed: Measured in seconds per epoch and number of epochs required to reach stable validation performance.
- (4) Model Complexity: Estimated by the number of trainable parameters.

These metrics offer a comprehensive view of model behavior beyond raw prediction accuracy, particularly important in clinical settings where false negatives can have serious consequences.

4.2 Comparative evaluation

Two quantum-enhanced approaches, VQCs and Quantum Kernel Methods, were compared against classical baselines: Multilayer Perceptron (MLP) and CNN with similar depth and parameter count.

The quantitative results in Table 1 reveal distinct performance patterns. Most notably, the Quantum Kernel SVM achieved the highest average precision (89.1%) and F1-score (88.5%) among all models, while the classical CNN attained the highest raw accuracy (89.1%). A key behavioral observation is that the Quantum Kernel SVM consistently demonstrated superior performance in correctly identifying minority classes, such as subtle hemorrhages and small tumors, where other models showed a higher rate of false negatives. This observed behavior could be explained by the underlying quantum mechanical principle of the method. The quantum feature map embeds classical data into a high-dimensional Hilbert space, where complex, non-linearly separable patterns may become more easily distinguishable. This enhanced separability directly translates to the higher precision and robust F1-score observed. These findings align with and are justified by the growing body of literature on quantum kernels; for instance, foundational work by Wei *et al.* [18] and recent applications in healthcare

Table 1. Average performance across datasets (5-fold cross validation)

Model type	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	Epochs to converge	Trainable parameters
MLP (Baseline)	87.6 ± 1.2	86.3 ± 1.4	85.9 ± 1.5	86.1 ± 1.3	42 ± 3	24,500
CNN (Baseline)	89.1 ± 0.9	88.2 ± 1.1	87.5 ± 1.2	87.8 ± 1.0	38 ± 2	35,200
VQC Model	86.8 ± 1.5	86.5 ± 1.6	86.1 ± 1.7	86.3 ± 1.5	26 ± 2	18
Quantum Kernel SVM	88.4 ± 0.8	89.1 ± 0.7	87.9 ± 1.0	88.5 ± 0.8	15 ± 1	8

Note(s): Values represent mean ± standard deviation across five-fold cross-validation on ChestX-ray14, Brain Tumor MRI, and CT Hemorrhage datasets

by Chow [8] similarly suggest that quantum feature maps can improve class separation for complex, high dimensional data, corroborating our empirical results.

Similarly, the VQC model demonstrated significantly faster convergence (26 epochs) compared to the classical MLP (42 epochs) and CNN (38 epochs), while using only 18 trainable parameters. This observation indicates that the hybrid quantum-classical optimization process can find a stable solution more efficiently. The likely cause for this is the expressive power of parameterized quantum circuits, which can represent complex functions with a relatively small number of parameters, potentially leading to a smoother optimization landscape. This efficiency in parameterization and convergence is a hypothesized advantage of variational quantum algorithms, as discussed in the work by Ranga *et al.* [7], and our results provide concrete evidence of this phenomenon in a medical imaging context.

To complement the quantitative results, the experimental observations revealed distinct behavior across imaging modalities. For example, the Quantum Kernel SVM achieved superior performance in detecting small or low contrast lesions (e.g. minor hemorrhages and early-stage tumors), while the VQC model demonstrated faster convergence on MRI datasets with limited sample sizes. These findings underline the adaptability of quantum enhanced models across diverse image characteristics, even when trained on simulated quantum backends.

The reported standard deviations indicate stable performance across cross-validation folds. While the absolute accuracy differences between quantum and classical models are modest, the quantum-enhanced models demonstrate substantial practical advantages in efficiency: parameter counts are reduced by over 99.9% (8 vs. 35,200 parameters) and convergence is 50–65% faster (15–26 vs. 38–42 epochs). These efficiency gains, rather than raw accuracy improvements, represent the primary value proposition for quantum-enhanced models in resource-constrained clinical environments.

Together with the variability metrics in Table 1, these results provide a comprehensive benchmark that reflects both model robustness and significant efficiency advantages of QML-enhanced architectures across distinct modalities and clinical conditions.

As illustrated in Figure 3, both quantum enhanced models, particularly the Quantum Kernel SVM, demonstrated significantly faster convergence compared to classical baselines. While the CNN and MLP models required 38 and 42 epochs respectively to reach stable validation performance, the VQC and Quantum Kernel models stabilized within 26 and 15 epochs. This accelerated convergence, coupled with competitive accuracy, highlights the potential of quantum models for efficient training, especially in resource constrained or time sensitive medical imaging workflows.

4.3 Simulated environment feasibility

All experiments were implemented in Google Colab, leveraging Qiskit and PennyLane simulators. The pipeline was tested across multiple sessions and consistently produced reproducible results within acceptable tolerances.

- (1) The reported per-epoch times (1.2–2.5 s for quantum models) are for simulations performed using the statevector simulator backend in Qiskit and PennyLane. This simulator efficiently computes the exact quantum state evolution for our 8-qubit circuits without sampling noise, which is feasible for circuits of this scale (8 qubits, depth ~20 gates). It is important to distinguish this from more computationally intensive simulation methods, such as density matrix simulation or shot-based sampling, which would indeed require significantly more time. The classical models (CNN/MLP) were run on the same Colab CPU/GPU environment for a direct comparison of computational overhead under these specific simulation conditions.




Figure 3. Validation accuracy vs. epochs for quantum and classical models

- (2) Memory usage was significantly lower for quantum models due to minimal parameter requirements.
- (3) Integration with classical frameworks (e.g. PyTorch via PennyLane) was seamless, enabling flexible experimentation without hardware constraints.

These findings confirm that simulated quantum environments are not only viable but also practical for early stage research in quantum enhanced medical imaging.

4.4 Interpretation and insights

The experimental results suggest that quantum enhanced learning offers distinct advantages in specific clinical contexts. It is important to note that the classical baselines were intentionally kept lightweight to match the scale of our quantum simulations. This controlled comparison highlights the efficiency gains of QML models rather than claiming superiority over heavily-optimized classical architectures.

- (1) High dimensional Feature Separation: Quantum kernels provide enhanced mapping of features into Hilbert space, leading to better separation of complex or subtle patterns, such as early stage tumors or microhemorrhages. To provide a quantitative basis for the discussed feature separation, we analyzed the per-class F1-scores from our experiments. The Quantum Kernel SVM showed a smaller standard deviation in F1-score across the 14 pathology classes in the ChestX-ray14 dataset ($\pm 3.2\%$) compared to the CNN baseline ($\pm 5.1\%$), indicating more consistent performance and suggesting the quantum feature map may help mitigate performance drops on challenging minority classes.
- (2) Efficient Optimization: VQCs demonstrate strong convergence properties with minimal hyperparameter tuning, making them suitable for limited-data scenarios.

- (3) **Reduced Model Complexity:** With orders of magnitude fewer parameters than classical models, quantum-enhanced architectures may support lightweight deployment in edge or resource-constrained clinical systems.
- (4) **Efficiency Advantages Quantified:** The reduction in parameters (>99.9%) and faster convergence (50–65%) provide concrete, quantitative evidence of efficiency gains, which is a primary advantage for near-term QML. The potential for improved interpretability through quantum state analysis remains an area for future research.

To provide a quantitative basis for the discussed feature separation, we analyzed the per-class F1-scores from our experiments. The Quantum Kernel SVM showed a smaller standard deviation in F1-score across the 14 pathology classes in the ChestX-ray14 dataset ($\pm 3.2\%$) compared to the CNN baseline ($\pm 5.1\%$), indicating more consistent performance and suggesting the quantum feature map may help mitigate performance drops on challenging minority classes. While not a direct visualization of feature space, this metric supports the observation of improved robustness in class separation.

While current limitations include slower simulation speeds and lack of real hardware integration, these challenges are expected to diminish as quantum computing platforms evolve.

These experimental findings can be contextualized within the emerging body of QML literature. It is important to clarify the scope of this comparison: this study does not claim that current hybrid QML models surpass the peak accuracy of state-of-the-art classical models like ResNet or Vision Transformers on well-established medical imaging benchmarks. The observed accuracy of 88.4% for our quantum-enhanced models is competitive with simpler classical baselines and with results reported in foundational hybrid QML studies [17]. The primary demonstrated value lies in the efficiency profile: achieving competitive diagnostic performance (within 1–2% of a classical CNN) while requiring orders of magnitude fewer parameters (8 vs. 35,200) and exhibiting significantly faster convergence (15 vs. 38 epochs). This efficiency makes them a compelling subject for further research, particularly for scenarios with stringent computational or data constraints where deploying massive models is impractical.

4.4.1 Clinical implications. These results open new possibilities for integrating quantum machine learning into clinical workflows, particularly in areas where:

- (1) Labeled data is scarce or imbalanced
- (2) Diagnostic speed and computational efficiency matter (e.g. emergency triage)
- (3) Interpretability is critical for clinician trust and regulatory compliance
- (4) Resource constraints limit deploying large-scale deep learning models (e.g. rural healthcare or mobile units).

We emphasize that the clinical applicability discussed here is based on the efficiency metrics (speed, parameter count) demonstrated in our simulation study. Full clinical integration would require validation on real hardware and within actual clinical workflows. As quantum hardware becomes more accessible, this study provides a foundational framework for future exploration of QML in scalable and accurate diagnostic modeling.

4.5 Limitations and future directions

All experiments in this study were conducted using idealized, noise-free quantum simulators (Qiskit and PennyLane). While this allows reproducible and accessible benchmarking of quantum-enhanced models, it does not capture the effects of decoherence, gate infidelity, or connectivity constraints present in real quantum hardware. Consequently, the observed

convergence and parameter efficiency advantages should be interpreted as simulation-based insights rather than guaranteed quantum speed-ups on physical devices.

Nevertheless, these results provide a proof-of-concept that hybrid quantum-classical models can achieve competitive performance with significantly fewer parameters, motivating further investigation on noisy intermediate-scale quantum (NISQ) hardware. Future work will focus on:

- (1) Benchmarking the proposed models on real quantum processors.
- (2) Investigating noise resilience and error mitigation strategies.
- (3) Scaling to higher qubit counts and more complex imaging tasks as hardware matures.

5. Conclusion and future work

This paper presented a hybrid quantum-classical framework for medical image analysis and demonstrated its utility for diagnostic classification tasks. Our primary conceptual contribution is this integrated workflow, which provides a standardized and accessible methodology for the biomedical AI community to benchmark quantum-enhanced models against classical counterparts. By leveraging hybrid quantum-classical models, implemented using variational quantum circuits and quantum kernel methods, this study demonstrated the analytical benefits of quantum learning in simulated environments.

The experiments confirmed that quantum-enhanced approaches could offer competitive performance compared to classical models, particularly in terms of model generalization, convergence efficiency, and robustness in small or imbalanced datasets. Importantly, the use of simulated quantum environments via tools like Qiskit and PennyLane, hosted on Google Colab, ensured that all models were accessible, reproducible, and suitable for early-stage research and experimentation. Our study demonstrates that hybrid quantum-classical models, even at current simulation scales, can match the performance of similarly-sized classical models while offering substantial efficiency benefits in parameter count and convergence speed. This positions QML as a promising direction for lightweight, efficient diagnostic tools in settings where deploying large-scale classical models is impractical.

For clinicians exploring this technology, we recommend starting with Quantum Kernel methods for classification using preprocessed imaging features, as they offer clear benefits with low implementation complexity. The significant reduction in model parameters (8 vs. 35,200) makes these approaches particularly suitable for deployment in resource constrained environments.

The results support the growing interest in the application of quantum computing within the healthcare AI domain, offering a novel computational perspective to enhance medical imaging workflows. The demonstrated methods are not intended to replace classical AI systems, but to complement them with additional support by a grant from the Research center of College of Computer and Information Sciences”, Deanship of Scientific capabilities that could become increasingly valuable as quantum hardware matures.

However, it is crucial to acknowledge the limitations of this study. Our work is based on simulations of quantum circuits, which operate under idealized, noise free conditions and are limited in qubit count. Real quantum hardware introduces decoherence and gate infidelity, which could impact model performance and remains a key challenge for practical application. Additionally, the computational overhead of classical simulation presents a current bottleneck. Furthermore, the classical baselines used for comparison (CNN, MLP), while standard for controlled efficiency benchmarking, and are not the most advanced models available (e.g. ResNet, Vision Transformers). Additionally, while we discuss potential advantages in feature separation, the study lacks direct visualizations of the quantum feature space (e.g. t-SNE plots) which would strengthen these claims. Besides, the reported training speeds are specific to

statevector simulation and do not reflect the times expected on actual quantum hardware or when using shot-based simulators with sampling noise, which are critical next steps for practical evaluation. Despite these limitations, which are inherent in most current NISQ-era QML research, our findings in a simulated environment provide a critical proof-of-concept and a foundation for future hardware deployment.

Future work will extend this study by experimenting with more complex imaging modalities, integrating multimodal clinical data, and benchmarking performance on real quantum hardware as it becomes more accessible. Additionally, further investigation into model interpretability, noise mitigation strategies, and domain adaptation will help bridge the gap between research and clinical application.

Author contributions

Soha Rawas: Conceptualization, Methodology, Software Implementation, Formal Analysis, Writing – Original Draft, Supervision. Duaa AlSaeed: Data Curation, Validation, Visualization, Writing – Review and Editing, Project Administration, Funding Acquisition.

Ethics approval

This study was exempt from ethics approval because it did not involve human or animal subjects. The data used in this study were publicly available and did not require informed consent from participants.

Use of generative AI tools

During the preparation of this work, the authors used generative AI tools (e.g. ChatGPT) to assist in refining the language, structuring the document, and improving clarity. The core ideas, scientific analysis, methodology design, implementation, and all critical thinking were solely the work of the authors.

References

1. Dhaini IN, El-Zaart A, Rawas S. Efficient and parallel medical image segmentation model (EPSM) based on brink-MCET using heterogeneous distributions. *Int J Online Biomed Eng.* 2025; 21(03): 3-164. doi: [10.3991/ijoe.v21i03.52189](https://doi.org/10.3991/ijoe.v21i03.52189).
2. Batool A, Byun Y-C. Brain tumor detection with integrating traditional and computational intelligence approaches across diverse imaging modalities-challenges and future directions. *Comput Biol Med.* 2024; 175: 108412. doi: [10.1016/j.combiomed.2024.108412](https://doi.org/10.1016/j.combiomed.2024.108412).
3. Rawas S, Dwinggo Samala A. EAFL: edge-assisted federated learning for real-time disease prediction using privacy-preserving AI. *Iran J Comput Sci.* 2025; 8(3): 1-11. doi: [10.1007/s42044-025-00251-x](https://doi.org/10.1007/s42044-025-00251-x).
4. Rawas S, Tafran C. Next-gen breast cancer diagnosis: iembc as an iomt-enabled cloud computing solution. *Multimedia Tools Appl.* 2024; 84(25): 1-28. doi: [10.1007/s11042-024-20358-w](https://doi.org/10.1007/s11042-024-20358-w).
5. Birjais R. Challenges and future directions for segmentation of medical images using deep learning models. In: *Deep Learning Applications in Medical Image Segmentation: Overview, Approaches, and Challenges*; 2025. p. 243-64.
6. Skiba R. *Quantum computing: fundamental principles of quantum computing systems*. After Midnight Publishing; 2025.
7. Ranga D, Rana A, Prajapat S, Kumar P, Kumar K, Vasilakos AV. Quantum machine learning: exploring the role of data encoding techniques, challenges, and future directions. *Math.* 2024; 12 (21): 3318. doi: [10.3390/math12213318](https://doi.org/10.3390/math12213318).
8. Chow JCL. Quantum computing and machine learning in medical decision-making: a comprehensive review. *Algorithms.* 2025; 18(3): 156. doi: [10.3390/a18030156](https://doi.org/10.3390/a18030156).
9. Chow JCL. Quantum computing in medicine. *Med Sci.* 2024; 12(4): 67. doi: [10.3390/medsci12040067](https://doi.org/10.3390/medsci12040067).

10. Sriraam N, Chinta B, Seshadri S, Suresh S. A comprehensive review of artificial intelligence-based algorithm towards fetal facial anomalies detection (2013–2024). *Artif Intelligence Rev.* 2025; 58(5): 1-49. doi: [10.1007/s10462-025-11160-7](https://doi.org/10.1007/s10462-025-11160-7).
11. Rawas S, Tafran C, AlSaeed D, Al-Ghreimil N. Transforming healthcare: AI-NLP fusion framework for precision decision-making and personalized care optimization in the era of IoMT. *Comput Mater Continua.* 2024; 81: 3-4601. doi: [10.32604/cmc.2024.055307](https://doi.org/10.32604/cmc.2024.055307).
12. Kumar RR, Priyadarshi R. Denoising and segmentation in medical image analysis: a comprehensive review on machine learning and deep learning approaches. *Multimedia Tools Appl.* 2024; 84(12): 1-59. doi: [10.1007/s11042-024-19313-6](https://doi.org/10.1007/s11042-024-19313-6).
13. Rawas S, Dwinggo Samala A. Revolutionizing brain tumor analysis: a fusion of ChatGPT and multi-modal CNN for unprecedented precision. *Int J Online Biomed Eng.* 2024; 20(08): 8-48. doi: [10.3991/ijoe.v20i08.47347](https://doi.org/10.3991/ijoe.v20i08.47347).
14. Rawas S. Transforming healthcare delivery: next-generation medication management in smart hospitals through IoMT and ML. *Discover Artif Intelligence.* 2024; 4(1): 31. doi: [10.1007/s44163-024-00128-1](https://doi.org/10.1007/s44163-024-00128-1).
15. Li X, Peng L, Wang YP, Zhang W. Open challenges and opportunities in federated foundation models towards biomedical healthcare. *BioData Min.* 2025; 18(1): 2. doi: [10.1186/s13040-024-00414-9](https://doi.org/10.1186/s13040-024-00414-9).
16. Ullah U, Garcia-Zapirain B. Quantum machine learning revolution in healthcare: a systematic review of emerging perspectives and applications. *IEEE Access.* 2024; 12: 11423-50. doi: [10.1109/access.2024.3353461](https://doi.org/10.1109/access.2024.3353461).
17. Arquam MD, Pathak V, Tripathi V, Panday R, Singh SK. Hybrid quantum-classical algorithms for neural network-based binary disease classification. In: *Asian Conference on Intelligent Information and Database Systems*; 2025. p. 337-48. Singapore: Springer Nature Singapore. doi: [10.1007/978-981-96-5887-9_25](https://doi.org/10.1007/978-981-96-5887-9_25).
18. Wei L, Liu H, Xu J, Shi L, Shan Z, Zhao B, Gao Y. Quantum machine learning in medical image analysis: a survey. *Neurocomputing.* 2023; 525: 42-53. doi: [10.1016/j.neucom.2023.01.049](https://doi.org/10.1016/j.neucom.2023.01.049).
19. Maheshwari D, Garcia-Zapirain B, Sierra-Sosa D. Quantum machine learning applications in the biomedical domain: a systematic review. *IEEE Access.* 2022; 10: 80463-84. doi: [10.1109/access.2022.3195044](https://doi.org/10.1109/access.2022.3195044).
20. Sonavane A, Jaiswar S, Mistry M, Aylani A, Hajoary D. Quantum machine learning models in healthcare: future trends and challenges in healthcare. In: *Quantum Computing for Healthcare Data*; 2025. p. 167-87. Elsevier. doi: [10.1016/B978-0-443-29297-2.00003-4](https://doi.org/10.1016/B978-0-443-29297-2.00003-4).
21. Wang X, Peng Y, Lu L, Lu Z, Bagheri M, Summers RM. Chestx-ray8: hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 21–26 July 2017, Honolulu, HI; 2017. p. 3462-71. doi: [10.1109/cvpr.2017.369](https://doi.org/10.1109/cvpr.2017.369).
22. Masoud Nickparvar Brain tumor MRI dataset. Kaggle; 2022. Available from: <https://www.kaggle.com/datasets/masoudnickparvar/brain-tumor-mri-dataset>
23. Helwan A, El-Fakhri G, Sasani H, Uzun Ozsahin D. Deep networks in identifying CT brain hemorrhage. *J Intell Fuzzy Syst.* 2018; 35(2): 2215-28. doi: [10.3233/jifs-172261](https://doi.org/10.3233/jifs-172261).
24. Lored R. Learn quantum computing with python and IBM quantum: write your own practical quantum programs with python. Birmingham: Packt Publishing; 2025.
25. Chakravarthi VS, Koteswar SR. Quantum system on chips. In: *SOC-Based Solutions in Emerging Application Domains*. Cham: Springer Nature Switzerland; 2025. p. 99-116.
26. Bisong E. Google colabatory. In: *Building Machine Learning and Deep Learning Models on Google Cloud Platform: A Comprehensive Guide for Beginners*. Berkeley, CA: Apress; 2019. p. 59-64.

Corresponding author

Duaa AlSaeed can be contacted at: dalsaeed@ksu.edu.sa

For instructions on how to order reprints of this article, please visit our website:

www.emeraldgroupublishing.com/licensing/reprints.htm

Or contact us for further details: permissions@emeraldinsight.com