Cureus Journal of Computer Science
https://doi.org/10.7759/s44389-025-00063-x

Cureus
Part of SPRINGER NATURE

Original Article

Quantum Artificial Intelligence for Medical Data: Comparative Evaluation of Quantum Machine Learning Models

Swathi Ganesan [1], Lakmali Karunarathne [1], Kavindu Karunarathne [1], Mishani Weerakoddi [1], Nalinda Somasiri [1]

1. *Computer and Data Science, York St John University, London, GBR*

Received: November 14, 2025 | Review began: December 29, 2025 | Review ended: April 14, 2026 | Published: May 05, 2026

Abstract

Background

Artificial Intelligence (AI) incorporation within healthcare applications has significantly helped in the improvement of healthcare diagnostics, patient forecasting, and clinical decision-making. The steady growing of medical data in terms of quantity, quality, and dimension poses more significant problems for traditional AI methods in terms of size, time, and computational efficiency. However, classical machine learning models often struggle with high-dimensional, imbalanced, and heterogeneous medical datasets, which limits their scalability and generalizability. The Quantum Artificial Intelligence (QAI) model, which combines quantum computing and machine learning in a complementary way, has emerged as one of the main trends to defeat these problems. This complementary approach is gaining more popularity, with the most advanced work being developed in this area, as it offers a fundamental explanation of current complexities and the underlying principles and their applications.

Methodology

This paper begins with the basic concepts of quantum computing and where it overlaps with AI. It then moves to a comparative approach, discussing various quantum machine learning algorithms such as quantum support vector machines, quantum neural networks, and variational quantum circuits. A comparative evaluation of quantum support vector machine, variational quantum classifier, quantum neural network, quantum K-means, and quantum Boltzmann machine is conducted using the Pima Indians Diabetes dataset to assess model behaviour and performance under standard metrics. Further detail is provided on the theoretical and practical implications of QAI in medical imaging, genomics, electronic health records, and drug discovery. Additionally, this paper examines platforms, tools, and quantum hardware used for QAI research, and discusses unique challenges currently faced in quantum technology.

Results

Unlike prior surveys that primarily provide conceptual overviews, this paper uniquely combines a comprehensive review with practical benchmarking of multiple quantum machine learning models on a benchmark medical dataset. This provides comparative insights into algorithmic performance under real-world constraints, particularly in relation to high-dimensional medical data, genomic data, electronic health records with large number of clinical attributes, radiology images, etc. In this study, dimensionality arises from multivariate clinical features encoded into quantum feature maps and the limitations of current quantum systems.

Conclusion

This paper aims to serve as a starting point for researchers and practitioners interested in the potential of quantum-enhanced AI to transform health data analytics. By identifying current limitations, evaluating the performance of emerging quantum machine learning approaches, and outlining gaps in the literature, the paper highlights future research directions necessary to advance QAI in healthcare. Overall, the findings outline both the potential and practical constraints of quantum models, offering a foundation for future work toward scalable, clinically applicable QAI solutions.

Categories: AI/ML-based decision support systems, Quantum Algorithms, Quantum Computing

Keywords: quantum artificial intelligence (qai), medical data analytics, quantum machine learning, healthcare applications, big data in medicine, hybrid quantum-classical algorithms, medical image analysis, genomic data processing

Introduction

The rapid transition to digital means in healthcare has created vast repositories of digital medical information, including records, images, genetic data, and wearables with data related to patient observation. The move to digitize purposefully created information has resulted in the largest collection of medical information to date [1]. This information is not conventional data, it can be multidimensional, sourced from various entities, and dynamically acquired, all of which present challenges in comprehension, cross-sectional study in real-time, and metrics. These characteristics make it increasingly difficult for classical analytical pipelines to manage scale, variability, and real-time decision-making requirements. Artificial Intelligence (AI) is starting to change treatment in this area, helping to identify diseases, generate prognosis implications, improve therapies, and ascertain medicines used [2]. The application of AI with medical data can facilitate healthcare systems in providing better choices, individualize healthcare for individuals, and improve patient care overall [3]. Despite these benefits, the growing complexity of medical datasets continues to expose structural and computational constraints within classical AI models.

Nevertheless, traditional AI models have challenges when addressing the complexities of larger medical datasets. Within this area are the curse of dimensionality, high computational costs, restrictions toward scalability, and challenges around extracting patterns from uncertain, incomplete, or unbalanced datasets [4]. In addition, deep learning models require a substantial amount of training data and access to significant computational resources, which may be a limit for applications involving rare diseases, or datasets where there are privacy restrictions [5]. As the amount and complexity of medical data increases, along with the constraints captured in this overview, the upsides of classical AI may be limited. This growing mismatch between data complexity and classical AI capability has led to exploration of alternative computational paradigms.

Quantum Artificial Intelligence (QAI) is a recently emerged intersection of quantum computing and AI that offers a potential way forward beyond current barriers to achieve better performance [6]. QAI leverages quantum superposition, quantum entanglement, and quantum parallelism, which may facilitate faster machine learning and improved and better pattern detection for large and complex datasets. Preliminary studies suggest QAI may also be of value in practical domains within medicine (e.g., medical image classification in radiology), genomic analysis, and drug discovery (e.g., medicinal chemistry), wherein speed of execution and ability to deal with complexity are likely to be critical factors [7,8]. Given these early indications, a structured and comprehensive examination of QAI's role in medical data analysis is required.

This paper aims to provide a comprehensive review of QAI applications in medical data analysis, examining existing approaches. The paper itself will focus on these mentioned aspects:

(1) To highlight the motivations and necessities for the QAI approach in healthcare

(2) To analyze the technical bottlenecks of classical AI for complex medical datasets

(3) To review the current state-of-the-art QAI methods in healthcare

(4) To explore remaining challenges and ongoing research opportunities in QAI

How to cite this article:
Ganesan S, Karunarathne L, Karunarathne K, et al. (May 05, 2026) Quantum Artificial Intelligence for Medical Data: Comparative Evaluation of Quantum Machine Learning Models. Cureus J Comput Sci 3 : es44389-025-00063-x. DOI https://doi.org/10.7759/s44389-025-00063-x

The paper will cover the developments of QAI algorithms, hardware, and hybrid quantum-classical models that have been translated across a variety of medical data modalities, representing the full landscape for researchers and practitioners.

Unlike existing surveys, which primarily concentrate on the theoretical landscape of QAI in healthcare, this study combines a comprehensive review with small-scale practical experiments across multiple quantum machine learning (QML) algorithms, including quantum support vector machine (QSVM), quantum neural networks (QNNs), quantum K-means (QK-Means), variational quantum classifier (VQC), and quantum Boltzmann machine (QBM), on the Pima Indians Diabetes Dataset (PIDD). By directly comparing algorithmic performance, class imbalance, and noise-related challenges under Noisy Intermediate-Scale Quantum (NISQ) conditions, this paper provides a distinctive contribution that bridges the gap between conceptual surveys and experimental validation. To contextualize this discussion, the following section outlines key quantum computing fundamentals relevant to QAI.

Background and related work

Quantum Computing Fundamentals

QAI is based on the principles of quantum computing methodology that uses the laws of quantum mechanics to process data at a level that classical systems cannot. A basic understanding of quantum computing is required to contextualize how QAI will solve the problems of medical data analytics [9]. This section summarizes quantum fundamentals and features large-scale, complex analytic data processing in healthcare. These fundamentals provide the groundwork for understanding why certain quantum operations may be advantageous when handling high-dimensional medical datasets.

1) Qubits: The building blocks of quantum computation begin with the concept of the quantum bit, or qubit, which underpins all subsequent QAI methods. The basis of quantum information is the quantum bit, or qubit, as shown in Figure *1*. Classical bits exist in a binary state of 0 or 1, but a qubit exists in superposition of both states [10].

2) Superposition: A single qubit can represent a number of states simultaneously in superposition. In turn, as the number of qubits increases, the computational space increases exponentially. The state of a qubit can exist in a superposition of states, thus allowing a quantum computer to evaluate a number of possibilities in parallel and potentially speed up solutions to certain problems [11]. Beyond superposition, another essential quantum phenomenon that contributes to QAI capabilities is entanglement.

3) Entanglement: This means a unique quantum correlation between qubits, meaning that one qubit's state is dependent on the other qubit's state, even when the qubits are separated from each other. Entanglement is a crucial element to quantum algorithms, quantum error correction, and quantum communication [12]. In QML settings, entanglement enables correlated feature representations that classical models cannot replicate.

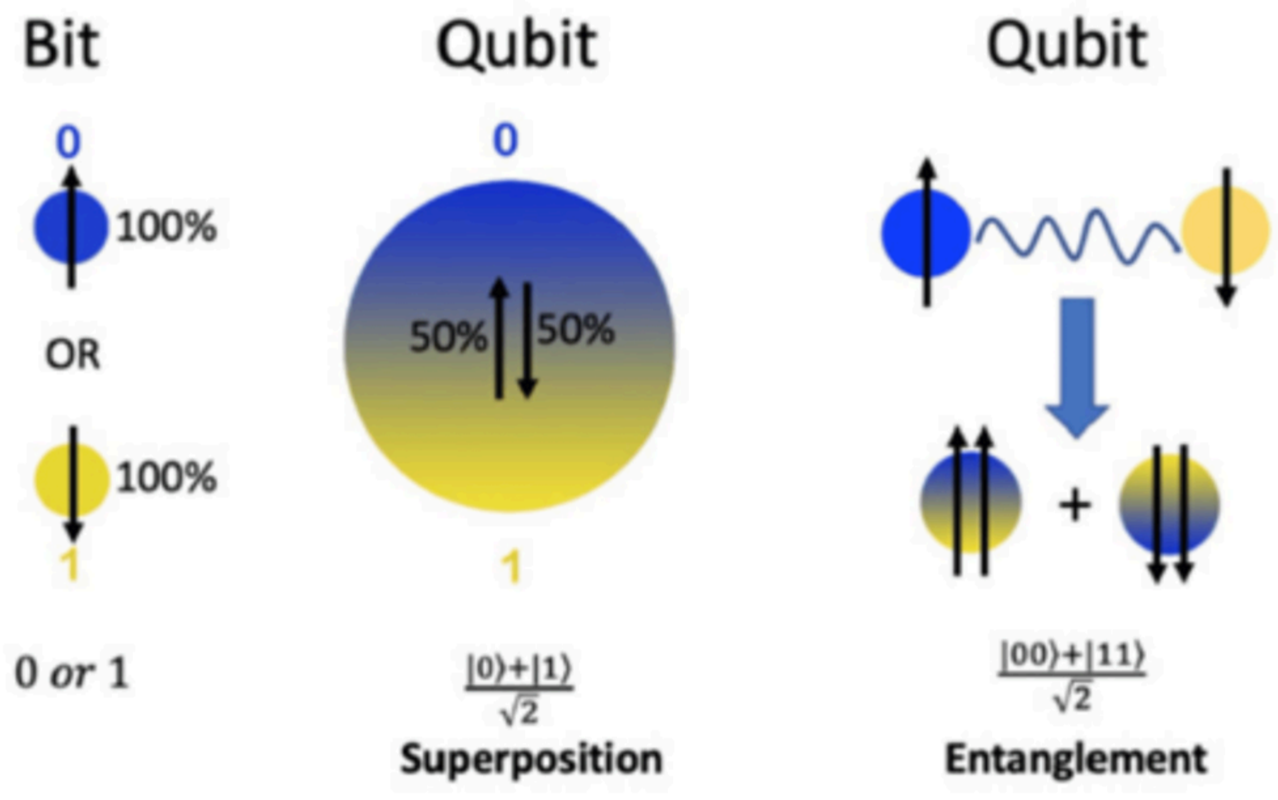

FIGURE 1: Qubit, Superposition, and Entanglement Illustration

[13]

4) Quantum gates and circuits: Quantum gates are basic processes used to manipulate a qubit through unitary transformation. With these foundational concepts established, quantum gates and circuits define how qubits are manipulated to perform useful computations. The concept can be accepted as the quantum gate, which is similar to logic gate in classical computing, but instead of working with a state space with discrete states and logical values, the quantum gate works with a continuous state space with probabilities. Quantum circuits are created by stringing together quantum gates to perform algorithms, the basic building block of quantum computing [14].

5) Differences from classical computing: Quantum and classical computing differ in how they handle information. This includes representation, storage, processing, and retrieval. In classical computing, information is encoded into discrete bits that reside in one of two states, either 0 or 1, while quantum computing uses qubits, which can appear to simultaneously be in multiple states, due to superposition, allowing exponential representation of information across increasingly numerous qubits [10], as shown in Figure *2*. These distinctions highlight why quantum systems may offer computational advantages for processing complex medical datasets that burden classical architectures. Classical processors tend to execute tasks sequentially with limited parallelism depending only on the number of hardware cores available, while due to superposition, quantum processors are able to use quantum parallelism to simultaneously explore many, many paths of computation. Classical systems are limited to either deterministic, or probabilistic correlations, while quantum systems have entanglement allowing for non-local correlations between qubits, thus enabling cooperative behavior that has no classical equivalent. Classical operations use deterministic logic gates (AND's, OR's, and NOT's), while quantum operations use quantum gates that are defined by their reversibility as unitary transformations and operate on continuous probability amplitudes, allowing for operations not possible in a classical system [15]. Furthermore, classical bits can be measured to reveal their value is stored directly, while quantum measurements effectively collapse a qubit's

state condition to one of its basis states with some probability, resulting in qubits requiring multiple runs on a computation to reliably obtain outcomes. With these quantum principles outlined, the next step is to understand how AI currently operates in healthcare before examining how quantum approaches may address its limitations.

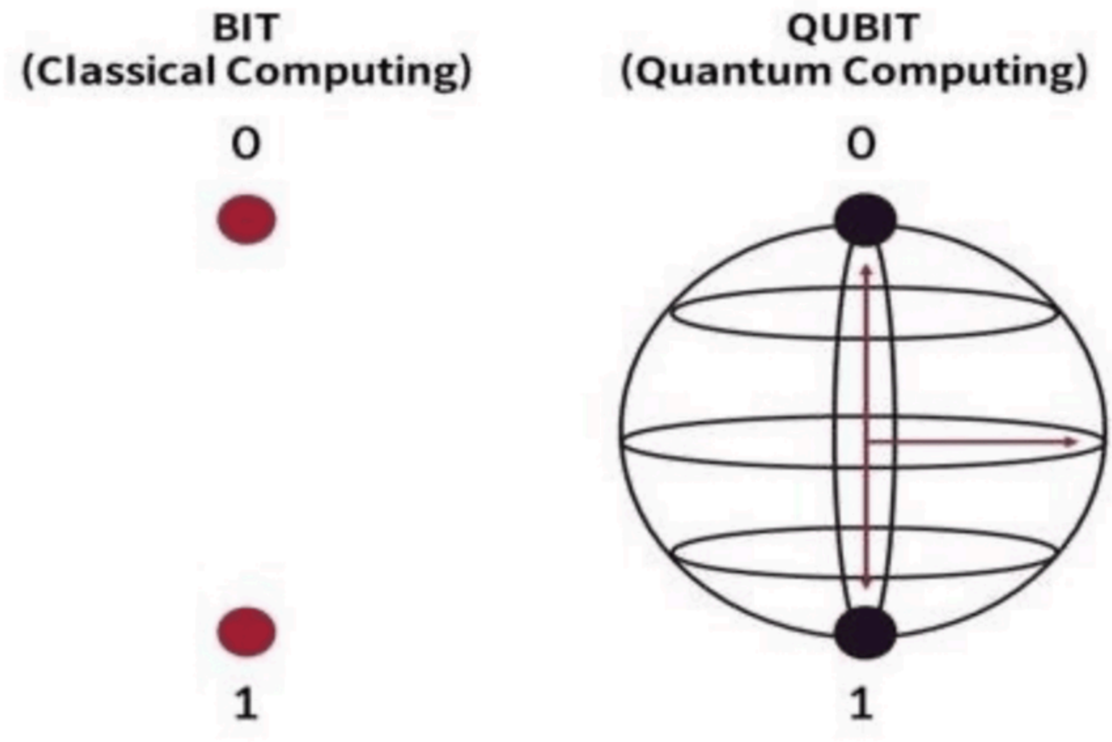

FIGURE 2: Bit vs Qubit Representation

[16]

AI in Healthcare

AI is a disruptive aspect of modern healthcare, providing new capabilities to analyze new and complex medical data. Conventional methods include machine learning and deep learning. These are being used more and more in clinical diagnosis, medical imaging, electronic health record (EHR) study, and genomic data work. Disease diagnosis, medical image interpretation, drug discovery, personalised medicine, robotic-assisted procedures, clinical trials, EHR analytics, and outbreak prediction are just a few of the areas in which AI is used in healthcare, as shown in Figure *3*. The combination of these things can lead to care that is more precise, efficient, and suited to the individual. AI models, like support vector machines (SVMs), decision trees, and neural networks, have also been used in diagnosis to sort diseases using clinical data, whether it is structured or not. Together, these capabilities offer more precise, efficient, and personalized care. However, the widening scope and complexity of healthcare data continue to expose critical weaknesses in classical AI pipelines. AI models, including SVMs, decision trees, and neural networks, have also been used in diagnostics to classify diseases from either structured or unstructured clinical data. As an example of a clinical application of predictive models, AI models can help clinicians make sense of clinical data through risk stratification, proactive disease detection, and predicting prognostic outcomes to produce more evidence-informed data-driven clinical decisions.

In the field of medical imaging, deep learning approaches, especially convolutional neural networks (CNNs), have proven to be state-of-the-art in tumor detection, organ segmentation, and radiologic pattern recognition [17]. Yet the success of these models is heavily dependent on data volume, computational capacity, and model interpretability. These methods have outperformed traditional image processing algorithms in speed and accuracy, enabling real-time decision support in radiology and pathology workflows. When analyzing EHRs, natural language processing and other machine learning methods are used in large-scale patient records to characterize meaningful patterns from unstructured notes, lab results, and prior histories. The EHR patterns can assist in predicting readmission events, identifying adverse drug events, and improving hospital resource allocation. In genomics, AI speeds up the discovery of genetic variants associated with disease risks and responses to therapies. Deep learning approaches can show complex gene-disease relationships and support precision medicine and personalized therapies [18].

Despite their potential, traditional AI methods have significant limitations when it comes to medical data. First, the high dimensionality and heterogeneous nature of datasets, such as imaging and genomics data, led to "computational bottlenecks," with immense training and storage costs associated with these high-dimensional ingrained processing methods. Second, prolonged training times associated with deep learning models made rapid deployment in emergent clinical settings, where time is critical, difficult, and result in long deployment cycles. Ultimately, considerable challenges remain with interpretable AI, where trust of clinicians with the use of non-interpretable mostly deep neural network models is lowered because the models worked as correlative "black boxes," leading to unforeseeable ethical dilemmas related to accountability. In essence, AI in healthcare has reached a turning point [18]. As these bottlenecks intensify, alternative computational models are increasingly explored to overcome constraints in scalability, speed, and interpretability. Its applications in healthcare will be significant and impactful, but all medical AI must address the bottlenecks above to have a usable, scalable, and ultimately trustworthy in the best medical settings. This need is perhaps why emerging paradigms like QAI are of interest since they may be able to overcome some of the challenges associated with the earlier sections by capitalizing on the benefits of quantum computational capacities as it pertains to classical AI. QAI presents a fundamentally different computational model that may address the shortcomings identified in classical frameworks.

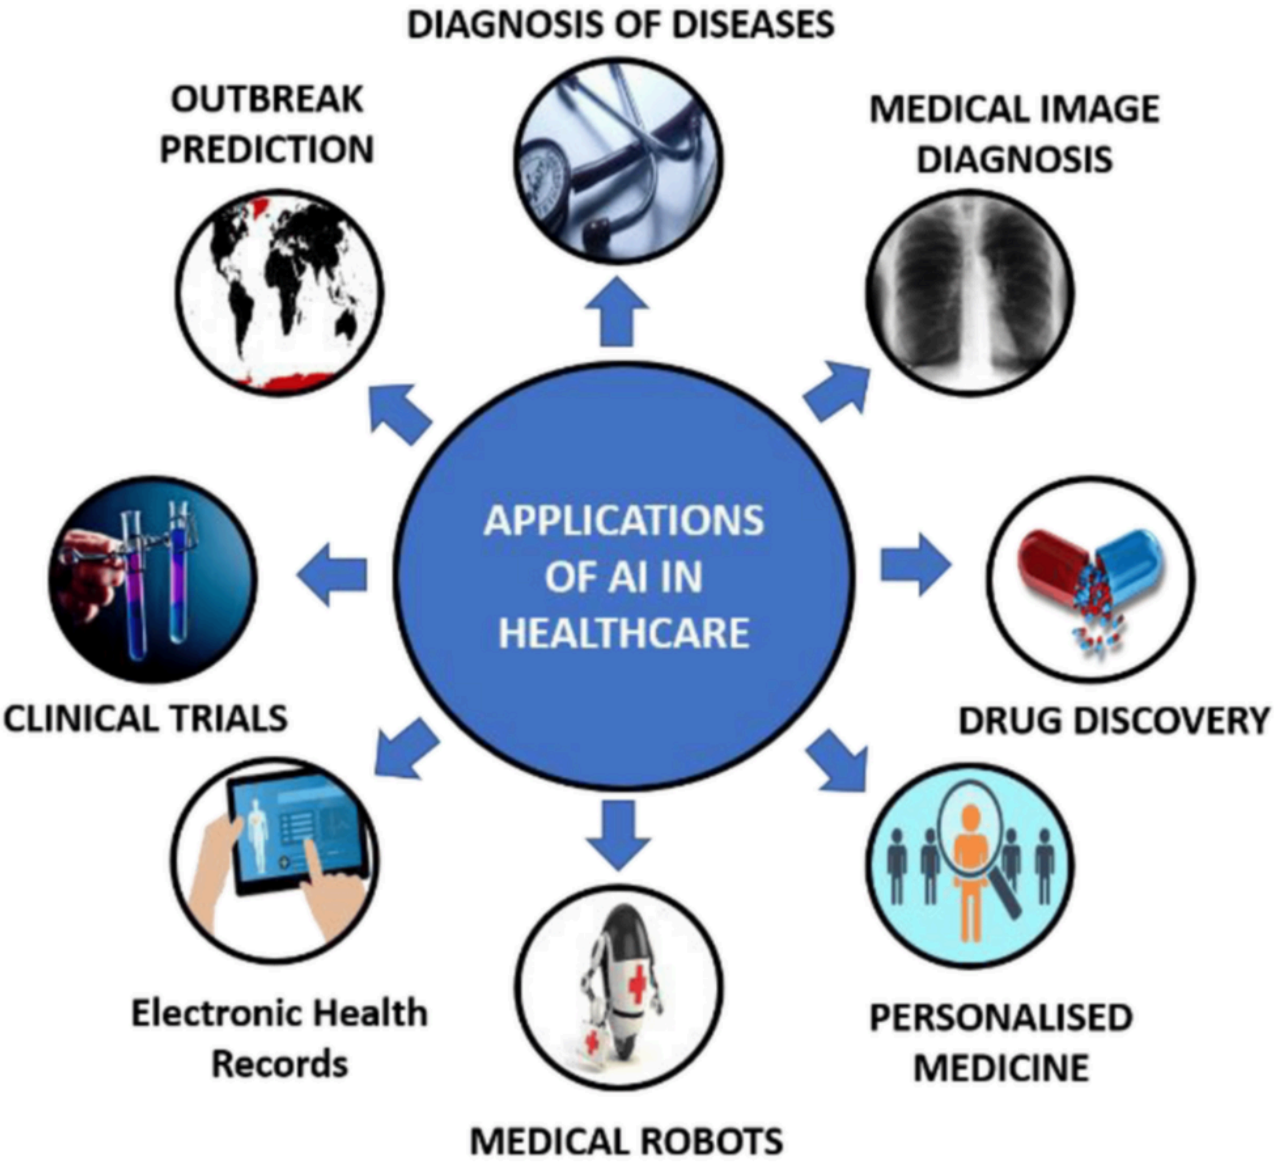

FIGURE 3: Applications of Artificial Intelligence in Healthcare

[19]

QAI and QML

QAI, also known as QML, is a new multidisciplinary field that combines AI methods with quantum computing methods. While classical AI extracts patterns from data using deterministic or purely statistical methods, QAI utilizes quantum methods such as superposition, entanglement, and quantum parallel evaluation to ultimately "pass" the computational barriers in classical methods [20]. Accordingly, QAI is inherently defined by performing learning models using quantum computations, enhancing or accelerating machine learning tasks to allow for learning from large, high-dimensional (or both large and high-dimensional) datasets that are impossible or unsolvable by classical systems. QAI is clearly useful in cases like medical data, heterogeneous and complex information that extends beyond the boundaries of classical AI systems. To understand how QAI models are implemented in practice, it is important to distinguish between the main categories of QML approaches.

Paradigms related to QML consist of three categories.

1) Quantum-inspired classical algorithms: Classical algorithms with quantum computing principles but require no quantum hardware.

2) Hybrid quantum-classical models: Algorithms that utilize quantum subroutines and classical workflows commonly in near-term quantum devices when they do certain tasks such as feature mapping or optimization.

3) Fully quantum models: Machine learning algorithms that work purely on quantum states (superposition) and quantum circuits with the development of future fault-tolerant devices.

Each of these models can collectively demonstrate a distribution from where we are today on a roadmap to long-term quantum-native AI systems.

Variational quantum algorithms (VQAs): VQAs are a fundamental type of near-term QML, since they are well suited for noisy, intermediate-scale quantum (NISQ) devices. A VQA is usually defined as a parameterized quantum circuit that has its parameters iteratively optimized by a classical optimizer to minimize a cost function [20]. VQAs can be defined to do clustering, classification, and regression, exploiting the quantum-induced benefits of new search and optimization landscapes. In medical domains, VQAs could not only help in drug discovery but also in personalized treatment predictions by balancing computational efficiency and the limitations of current quantum hardware.

Quantum-enhanced feature spaces: One of the most exciting QAI applications is the idea of quantum-enhanced feature spaces. This revolves around the belief that embedding classical data through quantum kernels into high-dimensional Hilbert space can lead to more efficient (in terms of data separation and concentration) separations (with QML versus classical kernel methods) in complex datasets. There are many applications of QML in industries like healthcare. For example, the high-dimensional feature maps in QML, or quantum-enhanced feature maps, may lead to the revealing of subtle patterns hidden in noisy and high-dimensional medical imaging or genomic data. In healthcare, quantum-enhanced feature maps may lead to significant leaps in predictive accuracy for identifying clinical diagnoses and biomarkers likely beyond the capabilities of classical AI. QAI/QML offers new paradigms beyond classical AI, and, hopefully, this will offer exponential speedups, richer representations of the data, and better scalability. We are at the inception of developing potential QAI/QML applications that are useful, as classical and quantum learning models are hybridizing to increase new applications in areas that are becoming more data-intensive like healthcare.

Literature Review on QAI in Healthcare

AI and quantum computing have recently gained attention as approaches to medical data analysis challenges. Recent research has explored the use of quantum methods for addressing medical data problems, building on the foundational ideas presented before. This section examines relevant studies in this area, with attention to their methods, applications, and limitations for future studies.

The potential impact of quantum computing and AI on healthcare analytics is a topic of great interest. QML seems like the next big thing for expediting medical judgements when dealing with complex clinical data, such as high-dimensional, noisy, or small datasets. However, it is evident from the literature that QML is still being established in the healthcare industry. After examining QML's role in digital health, it becomes clear that there is currently no conclusive evidence that QML performs significantly better than traditional machine learning. The majority of reported improvements are small, only appear in specific datasets, and typically occur in ideal simulation setups rather than on actual quantum hardware [21]. Although quantum computing shows promise for applications such as drug development, genomics, and diagnostic analytics, its true value currently lies in its integration with traditional systems. Simply put, fully quantum models are not yet available [22]. It is evident from all these findings that although quantum computers are frequently discussed as being superior, those assertions are still merely conjecture.

Another area of study focuses on specific QML algorithms, such as QSVMs, QNNs, and VQCs, and evaluates them using medical data. Benchmarking studies have have been conducted to compare these quantum models with standard classical ones. One empirical evaluation of QSVMs on a number of clinical datasets found that, in certain situations, QSVMs performed on par with or slightly better than classical SVMs, particularly when sample sizes were small but the data were high-dimensional [23]. However, the edge was somewhat slight and did not consistently hold; performance frequently declined when realistic noise was added. A head-to-head comparison between QSVMs, QNNs, and traditional models found that QSVMs tended to perform better with structured tabular data, whereas QNNs encountered optimization problems such as barren plateaus, a common issue in QML [24]. However, both studies highlighted significant limitations. It was noted that the datasets were limited, the classical models received little hyperparameter adjustment, and most of the work depended on simulators rather than real quantum computers. These factors make it difficult to generalize the findings. Moving forward, more rigorous and standardized evaluation methods for QML are clearly needed.

QML is no longer limited to tabular clinical data; it is now being used by researchers in biomarker discovery, genetics, and medical imaging. For instance, hybrid quantum-classical neural networks have been explored as substitutes for conventional convolutional models. A hybrid CNN was developed by Houssein et al. [25] to categorize COVID-19 chest X-rays. They used a limited dataset from a single source, and the quantum component of their model was only a shallow transition layer, which may explain the seemingly strong accuracy. Therefore, it is difficult to determine whether the quantum component truly had a meaningful impact or whether most of the performance was driven by classical preprocessing. An evaluation classified brain tumor MRIs using HQNet, finding that hybrid models can occasionally perform better than standard CNNs in situations with limited data [26]. However, larger and more diverse datasets are still needed to validate these conclusions. In genomics, QML approaches have been used to identify biomarkers associated with metastasis in clear cell renal cell carcinoma. These findings suggest that quantum kernels can detect distinct gene-expression patterns; however, their analysis has the same issue as many genomics studies using QML: a large number of variables with very few samples. Additionally, the lack of comparison with leading classical models such as gradient boosting or radial basis function SVM raises concerns about potential overfitting [27].

These literature studies reveal several recurring constraints and challenges. First, because existing NISQ hardware is plagued by decoherence, gate failures, and limited qubit connectivity, an overreliance on noiseless simulators leads to irrational assumptions about real-device performance. Second, small or biased datasets, sometimes drawn from a single clinical site, are commonly used in healthcare QML investigations, limiting statistical significance and generalizability. Third, the overhead of quantum data encoding is often overlooked when making claims about computational complexity. While methods like amplitude encoding require deep circuits that are unsuitable for noisy hardware, angle encoding limits information richness and may compromise theoretical advantages. Fourth, comparatively few QML studies include clinical validation, explainability analysis, or human-in-the-loop assessment, despite the critical need for transparency and trust in clinical settings.

Analysis indicates that QML will not outperform classical machine learning on large datasets until quantum hardware can correct errors [21]. Observations also suggest that quantum kernels are useful only in very specific situations. These ideas support the view that QAI may be more of a helper to classical machine learning rather than a substitute, at least when it is first used [28]. In general, current studies show that QAI is a growing but still young field. It presents interesting theoretical ideas but is constrained by limitations in both hardware and algorithms, and it often lacks real-world medical applications. For QAI to significantly impact healthcare analytics, more rigorous research is needed that considers hardware constraints and focuses on medical applications.

Relation to Prior Work and Broader Context

The experimental findings from this study complement several earlier works in classical and hybrid machine learning. Prior research has highlighted persistent challenges in traditional machine learning systems, particularly in handling high-dimensional, imbalanced, or noisy healthcare datasets [29]. Similarly, generative adversarial network-based approaches for synthetic data augmentation have shown promise in addressing data scarcity, though with limitations related to generative stability and clinical representativeness [30]. Additional work on optimized feature-selection frameworks has emphasized the importance of dimensionality reduction and model interpretability in improving predictive performance in medical applications. Collectively, these insights reinforce the need for continued exploration of quantum-enhanced and hybrid approaches in future healthcare analytics pipelines.

The research presented here further develops previous work in QAI by filling an important missing link between concept-level ideas of QML models and the actual testing of these models in real-world situations, specifically in the area of medical data. While previous research primarily focused on QML in simulation environments or projected future performance capabilities of QML models, this paper puts into practice many different QML algorithm implementations and provides a side-by-side analysis of their abilities (QSVM, QNN, QK-Means, VQC, and QBM) against a single benchmark for medical data (the Pima Indians Diabetes dataset). An evaluation of QML algorithms on NISQ devices demonstrates that quantum-enhanced solutions can address practical problems related to imbalanced classes, noise sensitivity, and data encoding in medical datasets. The novelty of this paper lies in its hybrid approach, as it offers a thorough analysis of the body of research on QAI while also presenting new empirical findings that validate algorithmic behavior, performance trade-offs, and practical constraints, thereby providing a more solid basis for the future clinical translation of QML technologies.

Materials And Methods

QML algorithms and methods

The QML algorithms try to combine the very principles of quantum computing with classical machine learning techniques for a better realization of data processing and pattern recognition. Based on quantum phenomena of superposition and entanglement, these algorithms have a great prospect of performing computationally complex operations at an exponential speed when compared to classical counterparts. QML investigates quantum models like QSVMs, QNNs, and VQCs to solve problems in optimization, classification, and data analysis, thus setting the path to realize breakthroughs in highly computationally demanding fields.

In order to find the studies that are relevant to QAI in the medical data area, a literature review was performed in a structured manner, which was semi-systematic in nature. The review was not only limited to the experimental applications of QML in healthcare but also included the conceptual framework.

In this study, PIDD [31] is used for the purpose of medical informatics and predictive modelling, for which it is a well-tested standard dataset. Originating from the National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK), the dataset contains 768 records of adult female patients of Pima Indian ancestry, all 21 years of age or older. Each record has eight independent variables: number of pregnancies, plasma glucose concentration, diastolic blood pressure, triceps skinfold thickness, two-hour serum insulin, body mass index, diabetes pedigree function, and age. In addition, there is one binary dependent variable (Outcome) indicating the presence (1) or absence (0) of diabetes. The dataset does not contain any explicit missing values but includes physiologically implausible zero entries in some attributes such as glucose, blood pressure, skin thickness, insulin, and body mass index, which are usually handled during the preprocessing step. The importance and well-structured nature of this dataset have led it to become a typical machine learning benchmark, especially for testing algorithms in the areas of early diabetes prediction and clinical decision support systems.

Quantum Support Vector Machine

QSVMs have calmed the realm of machine learning by merging quantum computation concepts with classical SVMs. This hybridization enriches the application of the quantum mechanics, enabling computational speed and increased accuracy in intricate classification scenarios. QSVMs employ quantum kernels and quantum-enhanced feature maps to act on data in a high dimension, something classical SVMs find cumbersome. The very potential of QSVMs can be demonstrated in areas like finance, medical diagnosis, and environmental science, emphasizing the ability to beat classical SVMs both in accuracy and in speed of computation. Some of the important considerations regarding QSVMs are given below [32].

The hybrid quantum classical SVM systems combine classic SVM concepts such as hyperplanes and kernel methods with those of quantum computing, which itself utilizes qubits, entanglement, and superposition. It allows the system to perform a job more rapidly and deal with data of a higher dimension. While the theory suggests that QSVMs would suffer the exponential speedup during computation and estimation of least squares SVM, they might have an advantage when it comes to processing complex datasets. Table *1* compares QSVM with classical SVM in terms of accuracy, scalability, and kernel methods. While QSVMs rely on quantum kernels, other QML approaches such as QNNs use parameterized circuits to learn more flexible representations [32].

Aspect	**Classical SVM**	**QSVM**
Accuracy	Mature, robust, well-optimized, accuracy depends on kernel choice and data.	May outperform classical kernels on certain tasks, but results are dataset-dependent and hardware-limited.
Scalability	Efficient for small to medium datasets, but scaling to very high dimensions or very large datasets is costly.	Theoretically scalable to higher dimensions via quantum states, but current quantum devices are limited in qubits and noisy.
Kernel Methods	Classical kernels	Quantum kernels generated from parameterized quantum circuits.

TABLE 1: Comparison of QSVM and Classical SVM

QSVM, Quantum Support Vector Machine; SVM, Support Vector Machine

Quantum Neural Network

QNNs represent a more prominent crossroad linking quantum computation with machine learning as they make use of quantum phenomena to gain greater computational power. Theoretically, processing that QNNs employ qubits and quantum operations may fill in the gaps where classical neural networks fall short of managing requirements for complex data or providing more speed in processing. The next section deals with some major features of QNNs [33].

A QNN is a hybrid machine learning model bringing together quantum computation and classical computation to exploit quantum advantages while still enjoying the classical optimization techniques. In a conventional QNN, one usually encodes classical data into a quantum state by angle encoding, amplitude encoding, basis encoding, etc. These states are then processed through a parameterized quantum circuit, which functions similarly to the hidden layers of a classical neural network, where quantum gates with tuneable parameters act on qubits [33].

Following the transformations, measurements are performed to yield classical outputs, which are compared against a cost function. A quantum classical training loop is then used, where a classical optimizer iteratively updates the parameters of the quantum gates. Gradient estimation methods such as the parameter-shift rule or more recent single-circuit approaches allow efficient training of these networks. To address optimization challenges such as barren plateaus, alternative designs like post-variational QNNs, which combine fixed quantum circuits with classical ensemble strategies, have also been proposed.

QK-Means Clustering

QK-Means clustering employs quantum mechanics to enhance the supremacy of the classical k-means for problems such as initialization dependence and noise resistance. Advancements in the field have brought forth many quantum variations that utilize, for instance, variational quantum eigen solvers or continuous-variable quantum computing while aiming for improved clustering performance and efficiency. The following sections discuss important developments in QK-Means clustering.

QK-Means is a quantum-inspired extension of the classical K-Means designed for more efficient clustering of large and high-dimensional datasets. Like the classical K-Means, QK-Means iteratively updates the position of the cluster centroids and/or assigns data points to the nearest cluster centroid. The main different is that QK-Means calculates distances and similarity measures using quantum circuits, whereby classical vectors are encoded into quantum states, mostly via amplitude or angle encoding [34]. A key technique used is the quantum inner product estimation (e.g., the swap test), which allows efficient evaluation of distances between high-dimensional data points in Hilbert space. After distances are measured, data points are assigned to clusters, and centroids are updated in a hybrid quantum classical loop. This approach offers potential polynomial or even exponential speedups for large-scale clustering tasks, though current implementations are limited by quantum hardware noise and qubit counts [34].

Variational Quantum Classifier

QML is a promising field that merges the representational power of machine learning with the computational advantages of quantum computing. Within this research area, the VQC has become one of the most widely examined algorithms for supervised learning on NISQ devices. A VQC uses parameterized quantum circuits that allow differentiation between various data classes. Essentially, quantum data encoding methods convert classical information into quantum states passed through the variational circuit, wherein VQCs employ features of a quantum system such as superposition and entanglement to draw decision boundaries that could otherwise have been difficult to represent by simple classical means [35].

So, unlike traditional quantum algorithms that require fault-tolerant quantum hardware, VQAs are devised as hybrid quantum classical algorithms. In these hybrid quantum-classical algorithms, quantum circuits serve for feature transformation and measurement, whereas classical optimization methods change the parameters of the circuit. Thus, they are particularly suitable for near-term devices. Yet, challenges such as barren plateaus and noise in the hardware still put in limits on its scalability. Nevertheless, VQC is still a primary contender to attest quantum advantage for practical classification tasks [35].

A VQC has an architecture composed of three steps: data encoding, variational processing, and measurement coupled with classical optimization. In this setup, classical data vectors are fed into quantum states via feature maps like angle or amplitude encoding, thereby representing data in the high-dimensional Hilbert space. The encoded states are then subjected to processing via the parameterized quantum circuit, which is composed of tuneable rotation gates and entangling gates and serves as the learnable layers of the classifier. The measurement of the quantum output state is then mapped to class labels. In the hybrid loop, the parameters of the circuit are updated by a classical optimizer that minimizes the value of some cost function, generally calculated by a gradient-based method such as the parameter-shift rule. By doing so, a VQC can portray complicated decision boundaries while being implementable on NISQ devices, although issues such as barren plateaus and quantum noise are still critical challenges.

Quantum Boltzmann Machine

The QBM is a quantum modification of the classical Boltzmann machine, a stochastic neural network used for generative modelling and unsupervised learning. While classical Boltzmann machines depend on probability distributions derived from Boltzmann distribution in statistical mechanics, QBMs extend the latter concept to quantum mechanics, where quantum Boltzmann distribution considers quantum properties such as superposition and entanglement. In addition, the training of QBMs can be carried out by hybrid quantum classical algorithms, where quantum processors prepare and measure states and classical computers update parameters [36].

The term QBM keeps being mentioned ever since the emergence of NISQ because quantum parallelism seems to model correlations that are classically unreachable. Opportunities for training, therefore, have arisen despite quantum noise, vanishing gradient problems, or simply the fact that preparing exact thermal states poses a big hurdle. Despite these drawbacks, QBMs remain of interest in QML [36].

In the context of healthcare, QBMs are particularly promising for generative tasks such as simulating synthetic patient data, modeling probabilistic dependencies among clinical features, and discovering hidden disease subtypes. Although current implementations are limited by noise and hardware constraints, these models offer potential pathways for addressing data scarcity and enhancing probabilistic reasoning in medical applications. The theoretical foundations and mechanisms of the key QML algorithms have now been outlined. The following section presents their practical evaluation on a benchmark medical dataset to examine how these models perform under real-world conditions.

Training, evaluation, and implementation framework

All QML techniques have been developed utilizing a mixed quantum-classical optimization framework. The measurement outcomes are generated by the execution of the parameterized quantum circuits during each training loop, and these outcomes are subsequently used to compute the classical cost function. Moreover, the updates of the model parameters are carried out by applying classical optimization algorithms iteratively, and this entire process is either continued until the convergence point is reached or until a pre-determined number of training iterations are done. To overcome the common problems associated with VQA training, like barren plateaus and noise susceptibility, shallow circuit designs and a small count of learnable parameters are continually employed in the trials.

The performance of the model is gauged through the application of standard supervised learning metrics like classification accuracy, precision, recall, F1-score, and confusion matrix analysis. Meanwhile, clustering performance for unsupervised learning models is evaluated through appropriate internal validation measures. The inclusion of classical machine learning models as baselines is intended merely to provide context for the performance of quantum models, and not to claim superiority or quantum advantage.

A NISQ-friendly quantum circuit simulator is employed for all the quantum models which ensures that the results are reproducible without depending on fault-tolerant quantum hardware. Data preprocessing, optimization, and evaluation of the classical parts are all performed using standard scientific libraries in a unified experimental environment.

The models of QML featured in this research paper are all developed using a mixed quantum-classical approach that is suited for NISQ devices. The classical parts are pre-processing, optimizing, and evaluating, while the quantum part is based on feature transformation and measurement using the circuit. The application of this methodology permits the simultaneous assessment of several QML paradigms on a common medical dataset in a consistent and reproducible manner.

Results And Discussion

Experimental setup and results

For the implementations, classical features were encoded into quantum states primarily using angle encoding, where each feature value is mapped to a rotation angle of a qubit. This approach was selected due to its efficiency on NISQ devices and compatibility with low-qubit experiments. Although amplitude encoding offers exponential data compression, it was avoided here because of circuit depth and noise sensitivity. Thus, the choice of angle encoding provided a practical balance between expressiveness and hardware feasibility.

Quantum Support Vector Machine

QSVM comes into existence as a quantum-empowered version of the classical SVM algorithm for data classification and pattern recognition. On the strength of the computing paradigm of quantum computation, and more specifically by applying quantum kernel estimation, the QSVM algorithm maps input data into a high-dimensional quantum feature space, thus accelerating the computation of complex relationships that are difficult to model classically. The QSVM can be advantageous for dealing with large datasets and improving classification accuracy; finance, bioinformatics, and image recognition are some areas where it may find application.

FIGURE 4: Quantum Support Vector Machine (QSVM) Accuracy

The QSVM was tested on a set of 154 samples, consisting of 100 samples from class 0 and 54 instances from class 1. The model achieved a gross accuracy of 64.9%, as shown in Figure 4, and an ROC-AUC of 0.625, implying some degree of discrimination capability between the two classes.

Referring to the confusion matrix in Table 2, it is seen that the QSVM correctly identified all 100 class 0 samples but could not classify any of the 54 class 1 samples. This implies that such a classifier would have achieved a perfect recall of 1.0 for class 0 but zero recall for class 1.

	Predicted 0	Predicted 1
Actual 0	100	0
Actual 1	54	0

TABLE 2: Quantum Support Vector Machine Confusion Matrix

The classification report shown in Table 3 stresses this very imbalance in the predictive performance. For class 0, the QSVM obtained a precision of 0.65 and an F1-score of 0.79, while for class 1, it has both precision and recall of 0.0, thus bringing down the F1-score to 0.0. The macro-averaged F1-score of 0.39 and the weighted F1-score of 0.51 show that the classifier is biased toward the majority class.

Class	Precision	Recall	F1-Score	Support
0	0.65	1.00	0.79	100
1	0.00	0.00	0.00	54

TABLE 3: Quantum Support Vector Machine Classification Report

How to cite this article:
Ganesan S, Karunarathne L, Karunarathne K, et al. (May 05, 2026) Quantum Artificial Intelligence for Medical Data: Comparative Evaluation of Quantum Machine Learning Models. Cureus J Comput Sci 3 : es44389-025-00063-x. DOI https://doi.org/10.7759/s44389-025-00063-x

The findings show that it was able to perfectly capture the structure of one big majority class from one to zero but completely failed to generalize to the other minority class 1. Such performance may be attributed partly to the class imbalance in the dataset and partly to the inability of the given quantum kernel to expressively separate minority-class samples. Moreover, the modest ROC-AUC score (0.625) flexibly permits the interpretation that the decision boundary by the model very ill-congeals for both classes in terms of robustification.

In practical terms, the QSVM, though suited to capture basic class structures, finds certain limitations when dealing with imbalanced data. Hence, one could:

(i) Test other quantum feature maps to obtain better class separations, or

(ii) Resample the data by oversampling the minority class or under sampling the majority class, or

(iii) Use a hybrid ensemble method whereby the decision bases of QSVM and classical models get balanced.

Resolving these limitations of QSVMs should pave the way to their applicability to real-world classification problems.

How to cite this article:
Ganesan S, Karunarathne L, Karunarathne K, et al. (May 05, 2026) Quantum Artificial Intelligence for Medical Data: Comparative Evaluation of Quantum Machine Learning Models. Cureus J Comput Sci 3 : es44389-025-00063-x. DOI https://doi.org/10.7759/s44389-025-00063-x

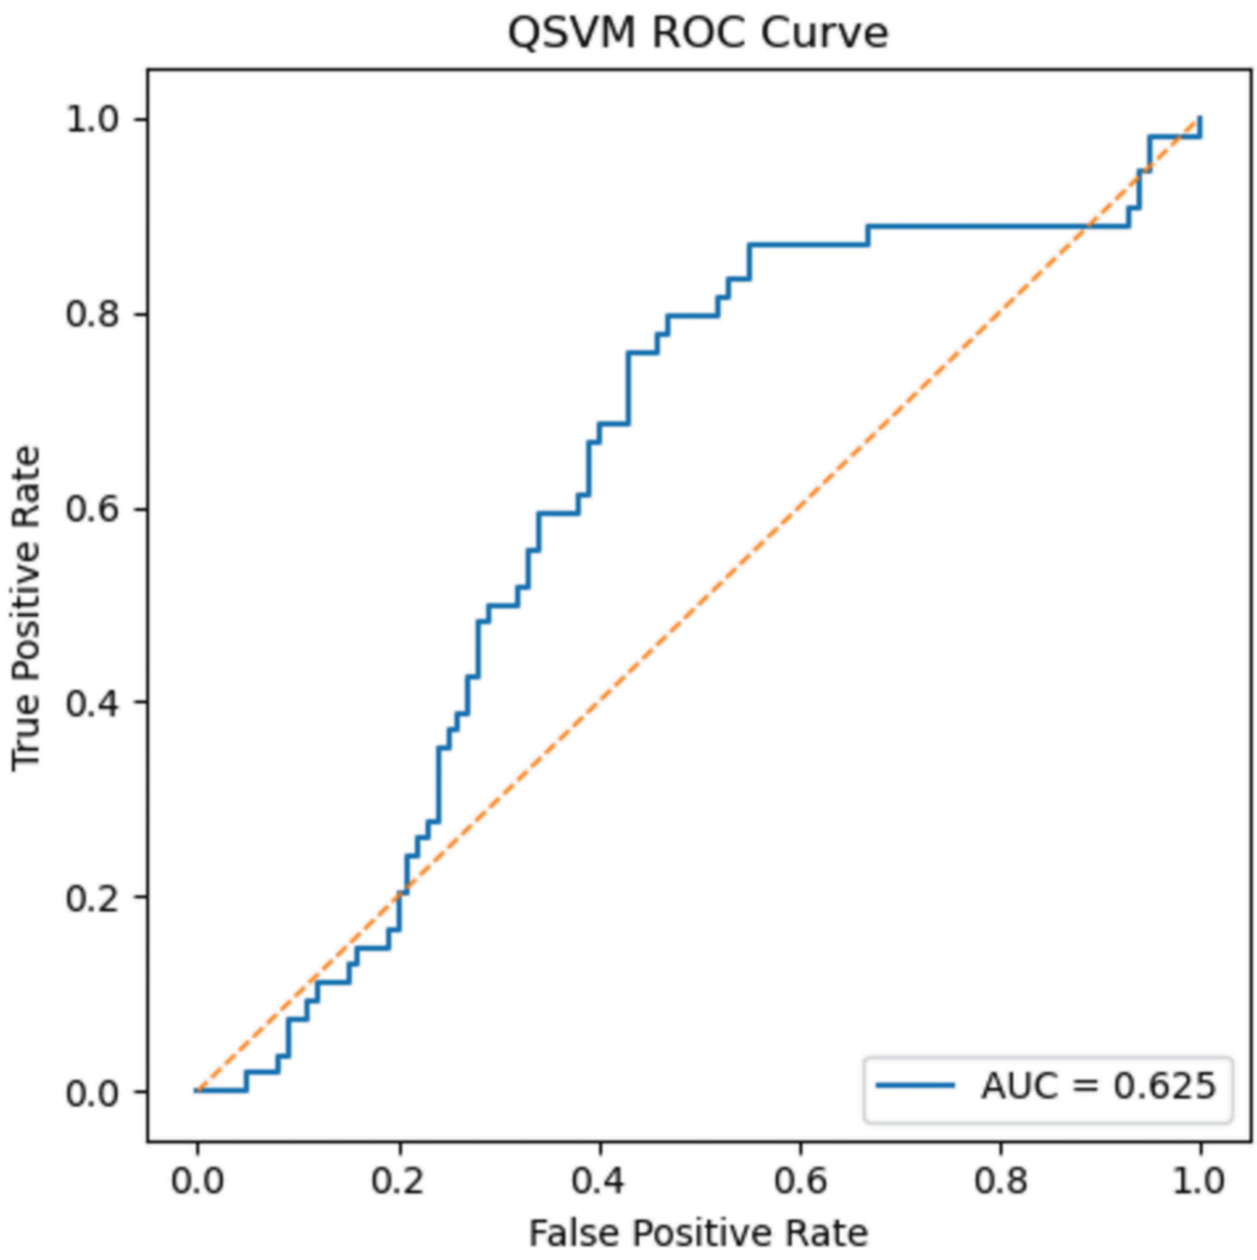

FIGURE 5: ROC Curve

AUC, Area Under the Curve; QSVM, Quantum Support Vector Machine; ROC, Receiver Operating Characteristic

The ROC curve for the QSVM, as shown in Figure 5, reinforces the weak discriminability power of the model between the two classes. The model obtained an AUC of 0.625, which is barely even slightly better than chance (AUC of 0.5). Hence, the QSVM-based modification provides only mild circumstantial evidence against random classification.

Some semblance of separation between true positives and false positives may have been observed in the curve; however, with such a low AUC, these findings align with the earlier confusion matrix shown in Figure 6. An obvious bias in favor of the majority class was displayed by the classifier. These results show that the quantum kernel can detect some patterns in the data but does not possess sufficient discriminative power for balanced classification.

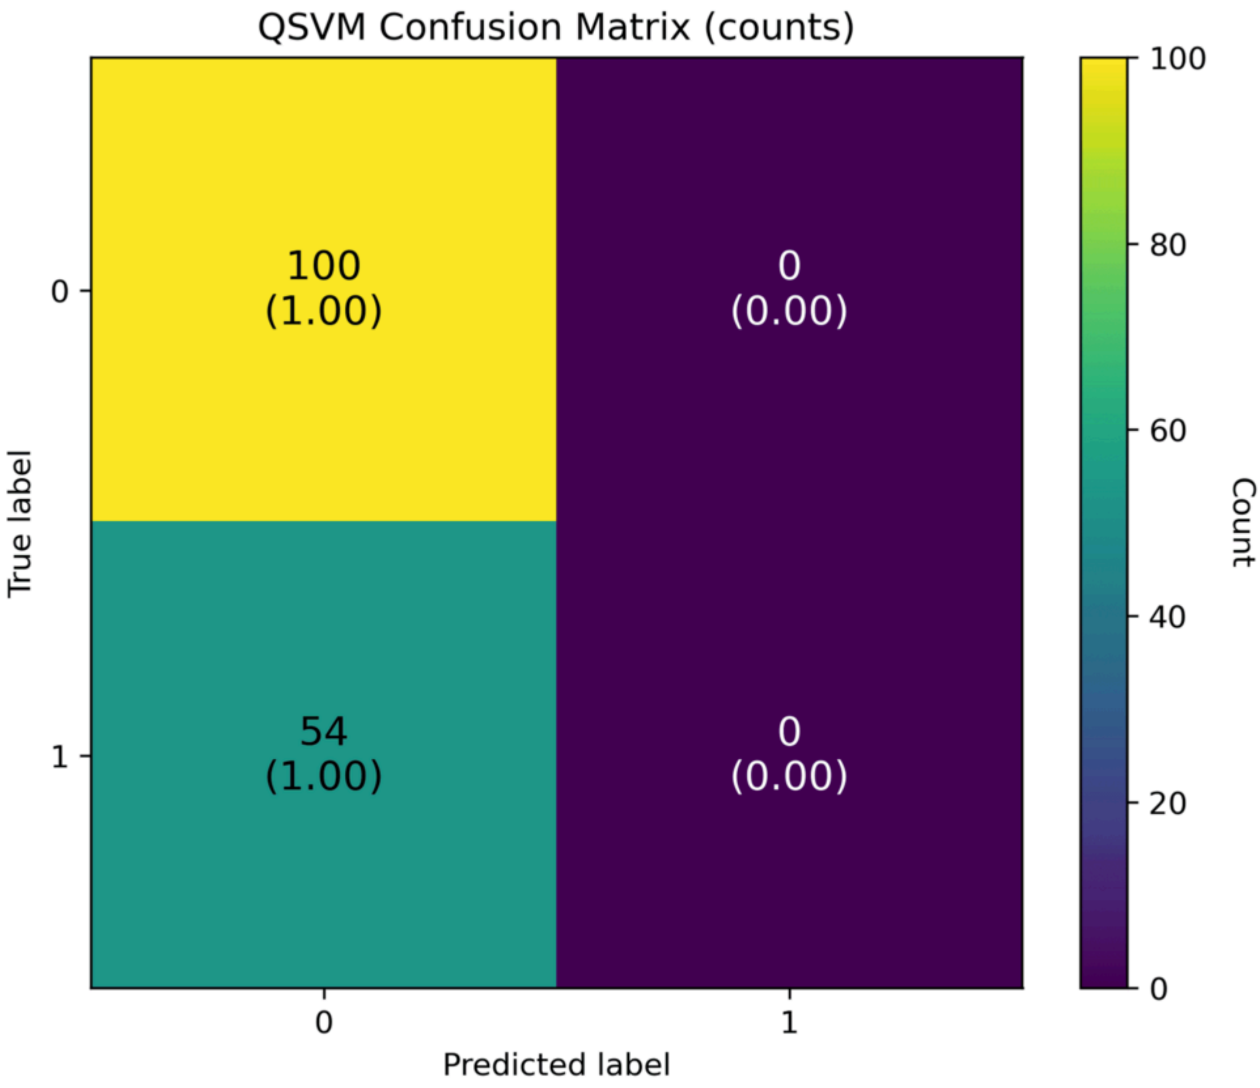

FIGURE 6: Quantum Support Vector Machine (QSVM) Confusion Matrix

Quantum Neural Network

A QNN is a quantum analogue of the classical neural networks used to aid learning and computation through quantum mechanics. By encoding data into a quantum state and then using the quantum gates as trainable parameters, a QNN can operate in complex pattern recognition through superposition and entanglement. The network would thus address such problems as optimization and processing of high-dimensional data, which have always been appropriately hard for classical systems. QNNs could also transcend quantum chemistry, materials science, and AI by providing faster and more potent forms of learning.

FIGURE 7: Quantum Neural Network (QNN) Accuracy

During training on 154 samples (100 from class 0, 54 from class 1), QNN showed an overall accuracy of 61.0%, as shown in Figure 7. The QSVM failed to classify any minority class samples, while the QNN was able to detect 19 class 1 instances, having a better treatment for imbalanced data problems. The results are presented in Table 4.

	Predict 0	Predict 1
Actual 0	75	25
Actual 1	35	19

TABLE 4: Quantum Neural Network Confusion Matrix

How to cite this article:
Ganesan S, Karunarathne L, Karunarathne K, et al. (May 05, 2026) Quantum Artificial Intelligence for Medical Data: Comparative Evaluation of Quantum Machine Learning Models. Cureus J Comput Sci 3 : es44389-025-00063-x. DOI https://doi.org/10.7759/s44389-025-00063-x

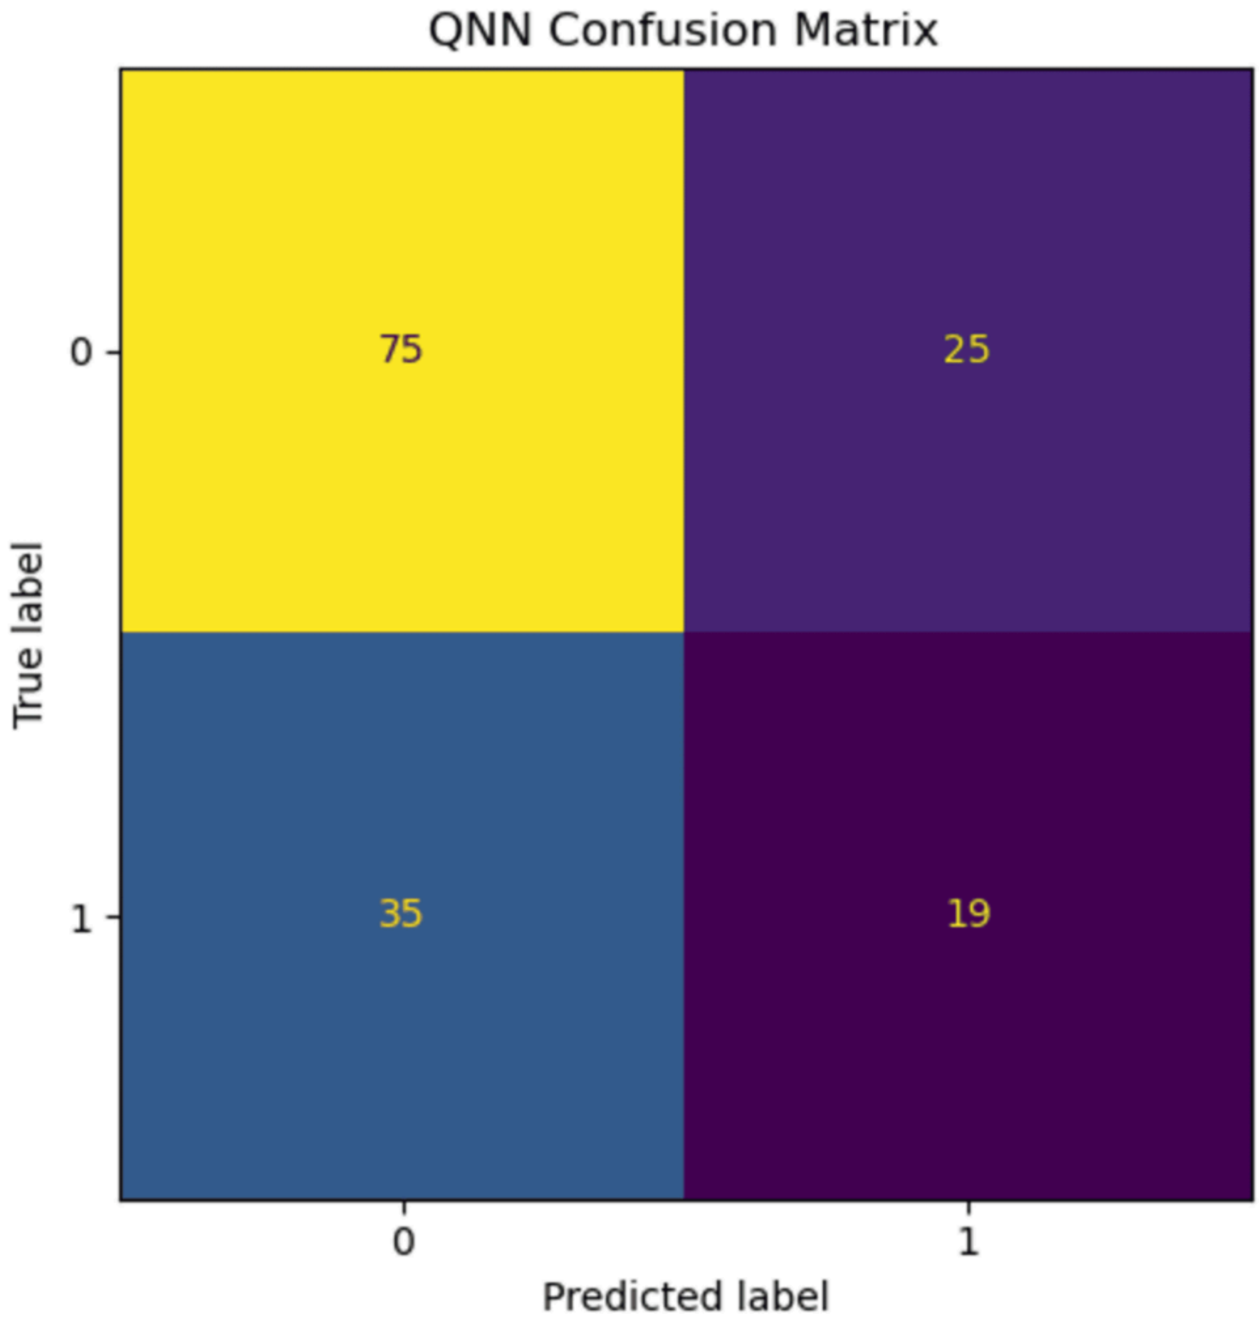

FIGURE 8: Quantum Neural Network (QNN) Confusion Matrix

From the QNN confusion matrix in Figure 8, while the model achieved correct classification in 75 instances out of 100 for class 0, it only correctly classified 19 out of 54 instances for class 1. While the QSVM failed to classify any of the minority-class samples, the QNN appears capable of partially capturing patterns from both majority and minority classes, albeit with limited precision. The model achieves a recall of 0.75 for class 0 and 0.35 for class 1, implying that although majority-class predictions are still stronger, the QNN manages to identify at least some minority instances. This results in a moderately balanced performance, with a macro-averaged F1-score of 0.55 and a weighted F1-score of 0.60, outperforming the QSVM's zero recall and F1-score for the minority class. While QNNs can handle class imbalance better than quantum kernel methods such as QSVM, the overall accuracy of 61% is still slightly on the lower side.

QK-Means Clustering

The QK-Means clustering algorithm attained an overall accuracy of 69.4%, emphasizing that the clusters in the quantum-enhanced feature space were separated better than by random assignment and fell just off what is in the range of accuracy obtained generally by classical clustering baseline. This level of accuracy suggests the quantum kernel embedding enabled the algorithm to capture some non-linear relationships inside the data and thus increase the consistency of cluster assignments.

This highlights the innate problem in the unsupervised technique, along with the sensitivity seen in the QK-Means to the quantum feature map's choice and the number of clusters, k. Furthermore, quantum noise and the depth of quantum circuits might have hindered the expressiveness of the algorithm, and this becomes especially relevant if the algorithm ran on NISQ hardware.

FIGURE 9: Cluster vs True Label Counts

The clustering of QK-Means, as shown in Figure 9, achieved an overall accuracy of 69.4%, validating the fact that the algorithm can discern meaningful structure in the dataset. Cluster C0 was dominated by class 0 samples: it assigned 333 samples correctly (83%), but the 17% (68 samples) from class 1 were assigned incorrectly. Cluster C1 was more evenly mixed, with 54% correct assignments of class 1 samples (200) against 46% incorrect assignments of class 0 samples (167). These outcomes suggest that QK-Means works better in identifying class 0 than class 1; it reflects the difficulty of dealing with overlapping feature distributions in unsupervised settings. Anyhow, the identification of 200 instances of class 1 shows that the quantum embedding was able to capture some non-trivial structure in the data that a strictly classical baseline might have had a tough time picking up on.

Variational Quantum Classifier

VQC achieved 71.7% accuracy in training and 73.4% accuracy in testing, as shown in Table 5, indicating that the model generalized slightly better on unseen data than on the training set. In effect, this showed that the parameterized quantum circuit learned meaningful decision boundaries without much overfitting, which for instance, quantum models commonly face due to their small or noisy nature.

In contrast with QSVM and QNN, VQC strikes a better balance between training and testing accuracy, giving credence to variational circuits in modelling complex patterns. The higher test accuracy, however, might indicate better generalization provided by the quantum feature space representation of VQC, despite the constraints posed by the limited number of qubits and gate noise.

Metric	Accuracy
Training Accuracy	0.717
Testing Accuracy	0.734

TABLE 5: Variational Quantum Classifier Accuracy

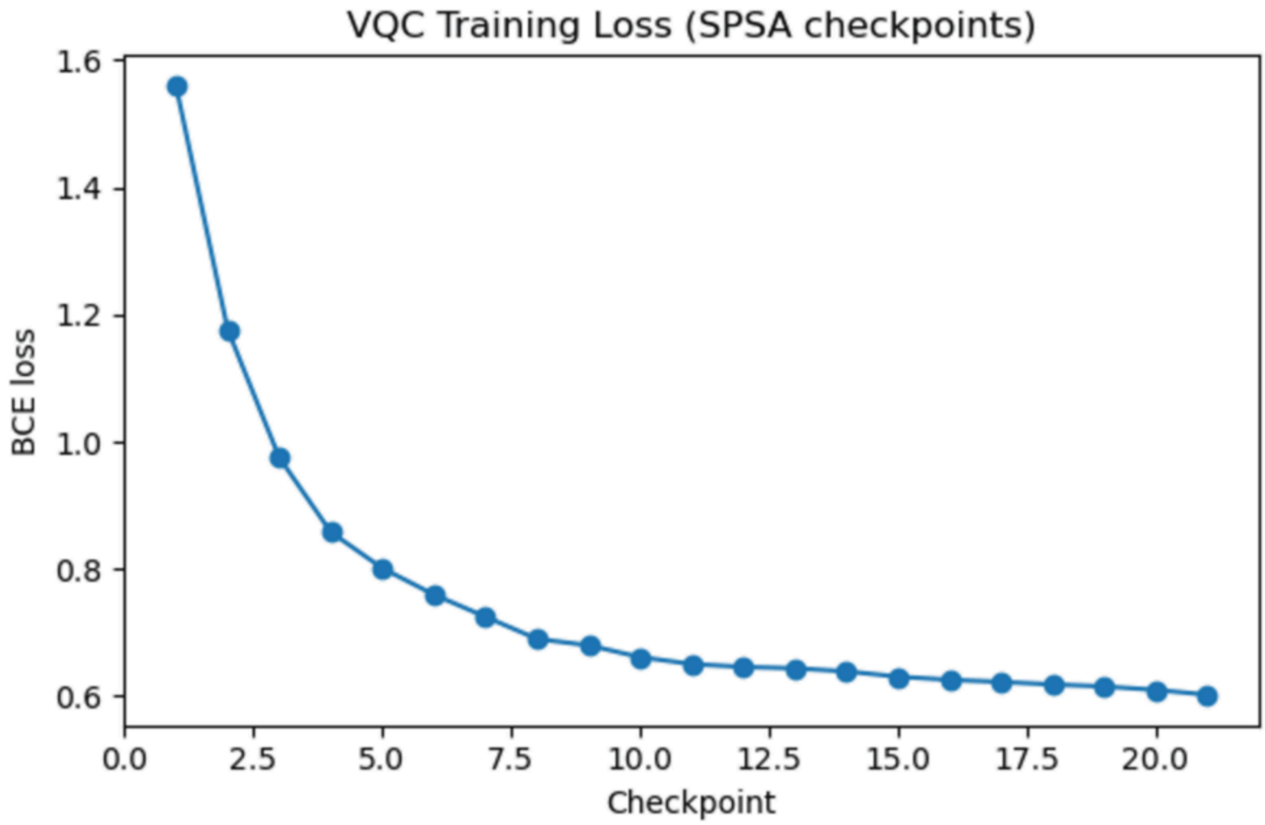

FIGURE 10: Variational Quantum Classifier (VQC) Training Loss

SPSA, Simultaneous Perturbation Stochastic Approximation

The VQC training loss, as shown in Figure 10, shows a clear descent along the optimization axis. The BCE loss with a starting value of around 1.55 steadily decreased until the final checkpoint, close to around 0.60. A steep drop was noted during the early few iterations, probably because the greedy optimizer quickly found the main data patterns, whereas the later curve with the uniform descent suggests the stable behavior of training. The nice and smooth downward trend echoes the fact that the Simultaneous Perturbation Stochastic Approximation (SPSA) optimizer was free to roam through the noisy and high-dimensional parameter space of quantum circuits. More importantly, the absence of large oscillations or divergences in that curve further suggests that the variational circuit ansatz and its chosen hyperparameters fit this classification task. Together with the 73.4% achieved test accuracy, the training curve validates that the VQC generalized well and did not in any way severely overfit.

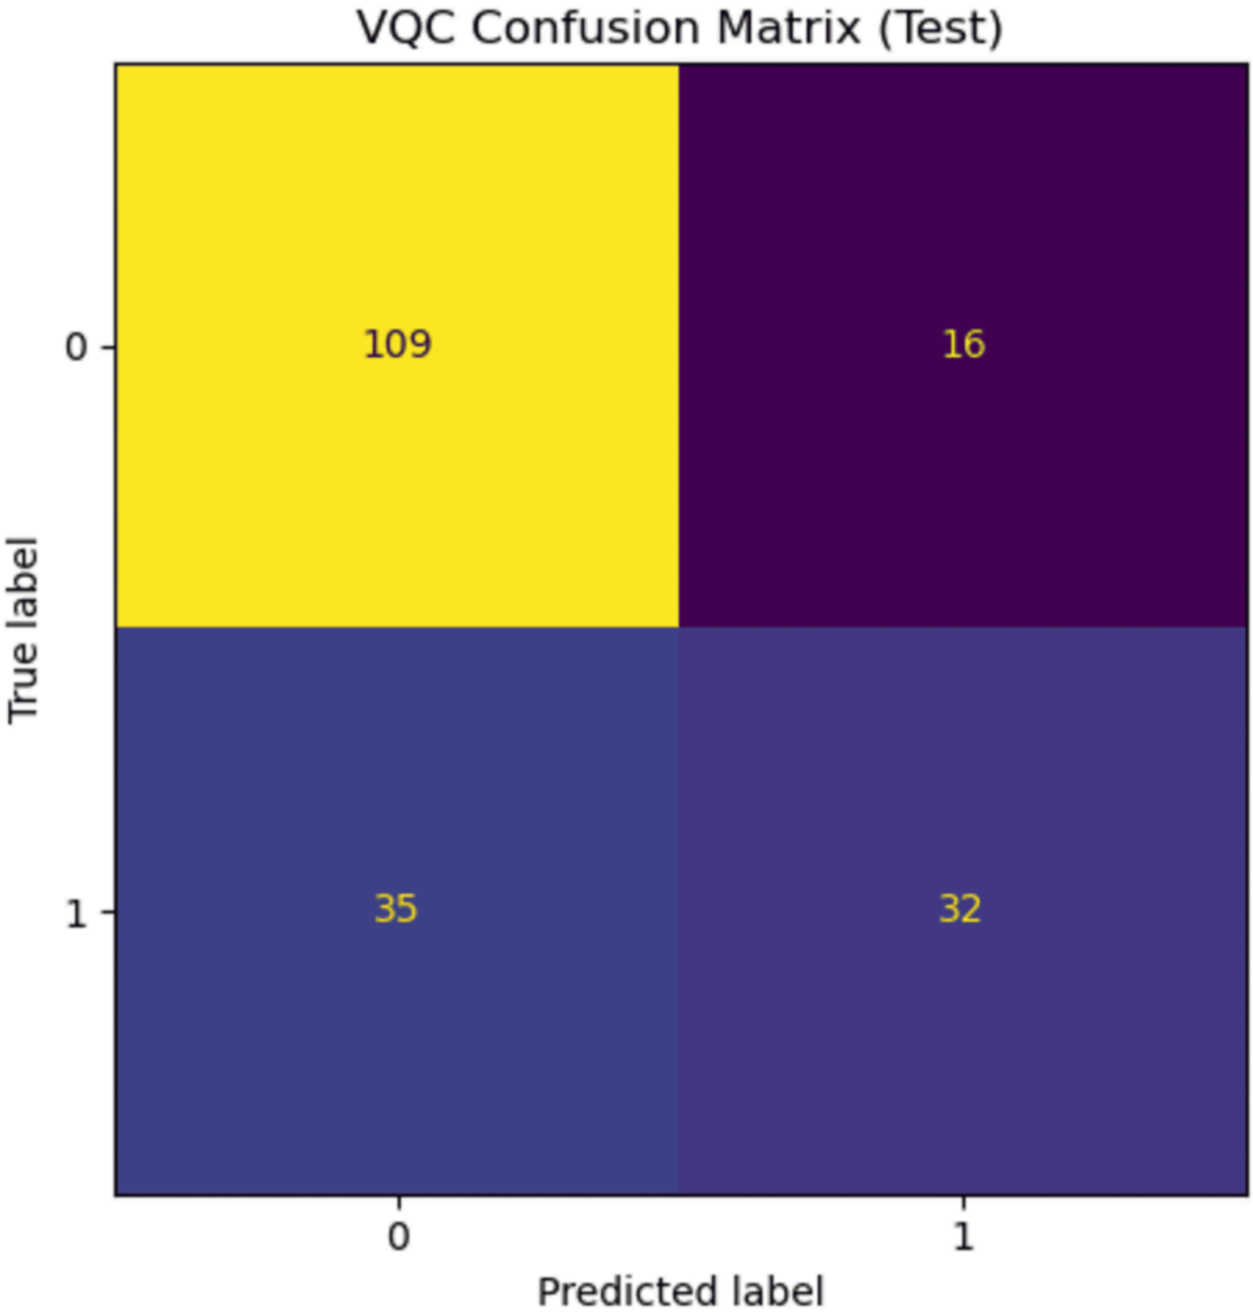

FIGURE 11: Variational Quantum Classifier (VQC) Confusion Matrix

In the test set confusion matrix for the VQC in Figure 11 , the model has correctly predicted 109 instances of class 0 out of 125, and 32 instances of class 1 out of 67. Faulty categorizations consist of 16 false negatives (class 0 identified as class 1) and 35 false positives (class 1 identified as class 0).

The results translate to a test accuracy of 73.4%, with somewhat excellent performance on the majority class (class 0 accuracy: 87.2%) and moderate on the minority class (class 1 accuracy: 47.8%). Compared to the QSVM and QNN, the VQC stays much more balanced between both classes and manages to extract some patterns from minority-class samples without sacrificing good performance on the majority-class samples.

Quantum Boltzmann Machine

The QBM received training on four features (f0, f1, f2, f3), and the generative potential of the model was tested by comparing the distributions formed on samples generated by the model with the empirical distribution of the data. The distribution distances had a total variation (TV) distance of 0.532 as shown in Table 6 and a KL divergence (data model) of 1.137, which means the QBM was able to capture some aspects of the data's structural patterns while still leaving a major gap between the true and model distributions.

Model-generated samples in Table 6 such as "0111," "1010," "0000," and "1111" reveal that the QBM persistently generated recurrent bit-strings that correspond to the high-probability patterns. The top states generated by the model were 0000 (22.7%), 1010 (18.8%), 0101 (15.9%), and 0111 (15.8%), which partially overlapped with the top empirical states such as 1111 (16.7%), 1011 (10.3%), and 0000 (9.9%). Such an overlap can be interpreted as the QBM having learned the low-energy states representative of the data at least partially but biased towards simpler states like 0000 and 1010.

In conclusion, the QBM succeeded in giving a rough outline of the true probability distribution, but some states were overrepresented when contrasted against the empirical distribution. The pretty high TV distance and KL divergence values offer good pointers where there is room for improvement: better Hamiltonian parameters and deeper circuit architectures or even better training paradigms. Even so, the QBM looks promising as a generative quantum model that can host complex probability landscapes.

How to cite this article:
Ganesan S, Karunarathne L, Karunarathne K, et al. (May 05, 2026) Quantum Artificial Intelligence for Medical Data: Comparative Evaluation of Quantum Machine Learning Models. Cureus J Comput Sci 3 : es44389-025-00063-x. DOI https://doi.org/10.7759/s44389-025-00063-x

Used Columns	['f0', 'f1', 'f2', 'f3']
Yotal Variation Distance	0.532
Kullback–Leibler Divergence (data ‖ model)	1.137
Model Samples (first 20)	0111, 1010, 0111, 0101, 0000, 1111, 0111, 0111, 0000, 1010, 1010, 1111, 0101, 0111, 1010, 1000, 0001, 0000, 0111, 0101
Top Empirical Distributions	1111 - 0.1667
	1011 - 0.1029
	0000 - 0.0990
	0111 - 0.0872
	1101 - 0.0677
Top Model Distributions	0000 - 0.2267
	1010 - 0.1883
	0101 - 0.1594
	0111 - 0.1576
	1000 - 0.1127

TABLE 6: Quantum Boltzmann Machine Results

The training curve shown in Figure *12* of the QBM reveals a clear convergence of the model with respect to squared Maximum Mean Discrepancy (MMD^2). The loss was high in the early iterations, near 0.64, but it drastically diminished to below 0.1 by checkpoint 8. It continued decreasing at a steady pace until stabilizing at around 0.05, indicating that the model distribution and empirical data distribution became better aligned.

The monotonic decrease in MMD^2 demonstrates that the QBM has learned to approximate the target probability distribution sufficiently well with stable convergence, without any instances of divergence or oscillatory behavior. This training curve, when supplemented with the sampling results (in which several high-probability empirical states overlapped with those of the model), offers quantitative validation of the QBM's ability to capture relevant structural patterns in the data.

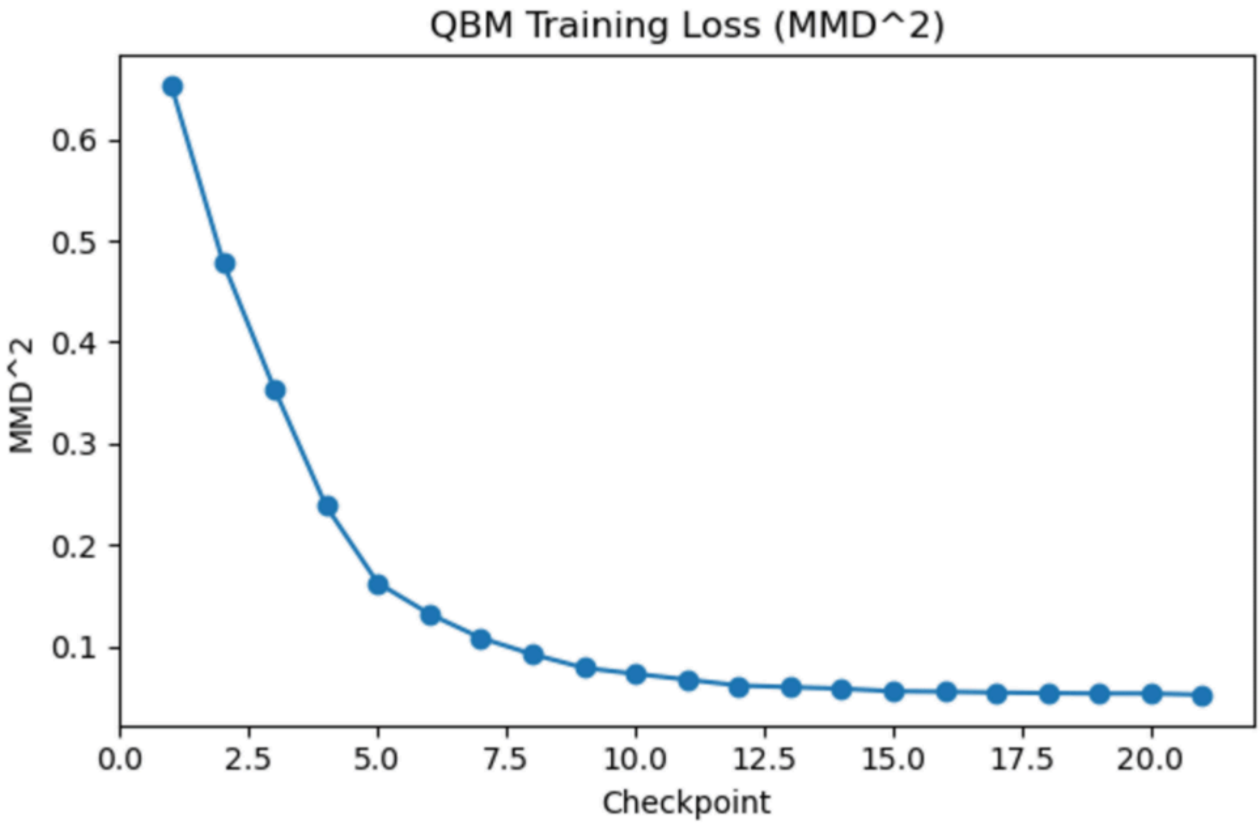

FIGURE 12: Quantum Boltzmann Machine (QBM) Training Loss

MMD, Maximum Mean Discrepancy

Model	Task Type	Accuracy	Strengths
QSVM	Supervised (Classification)	Accuracy = 0.649, ROC-AUC = 0.625	High accuracy on majority class (100/100 correct)
QNN	Supervised (Classification)	Accuracy = 0.610, Macro-F1 = 0.55	Balanced performance across classes, detected minority class (19/54 correct)
QK-Means	Unsupervised (Clustering)	Accuracy = 0.694	Captured non-trivial structure, reasonable cluster separation (PCA)
VQC	Supervised (Classification)	Train = 0.717, Test = 0.734	Best overall accuracy, balanced generalization, stable training loss curve, improved minority detection (32/67 correct)
QBM	Generative (Modelling)	TV = 0.532, KL = 1.137	Learned recurring patterns, captured several high-probability states.

TABLE 7: Quantum Machine Learning Model Comparison

QSVM, Quantum Support Vector Machine; QNN, Quantum Neural Network; QK-Means, Quantum K-Means; VQC, Variational Quantum Classifier; QBM, Quantum Boltzmann Machine; ROC, Operating Characteristic; AUC, Area Under the Curve; PCA, Principal Component Analysis

Overall, the comparative results shown in Table 7 highlight that different QML models excel in different aspects of medical data analysis. The QSVM achieved high accuracy on the majority class but failed on minority samples, reflecting its sensitivity to class imbalance. The QNN provided more balanced performance across classes, though with lower overall accuracy. The QK-Means clustering revealed non-trivial structure in the data with nearly 70% accuracy, indicating potential for unsupervised patient subgrouping. The VQC achieved the best overall performance with 73.4% test accuracy, showing strong generalization under NISQ constraints and making it the most promising supervised approach for near-term clinical tasks. Finally, the QBM captured key distribution patterns, underscoring its relevance for generative modeming in healthcare. Collectively, these findings demonstrate that while QML is not yet ready to surpass classical machine learning at scale, certain algorithms - especially variational methods - already provide valuable insights for medical data applications.

A comparative analysis of QML and classical machine learning approaches

Computation, Representation, and Complexity

Classical machine learning is an old and well-established discipline that has been heavily optimized and is reliably reproducible across CPU, GPU, or distributed computing clusters. Contrarily, QML poses constraints given its current-generation quantum hardware, which is NISQ devices. These systems exhibit limited numbers of qubits and coherence time while maintaining a reasonable number of errors, thus barring large-scale practical deployments. Consequently, present-day QML implementations are mostly run on quantum simulators, if not hybrid classical-quantum systems.

QML has theoretical advantages from a computation complexity perspective, which justifies its utterance. For example, quantum kernel methods have the potential to map classical data systematically into exponentially large Hilbert Space despite the circuit being polynomial size; while classical kernel expansions grow combinatorically with the feature dimension. A similar relation is true for certain algorithms, such as quantum principal component analysis, which have polylogarithmic complexity in the number of features as opposed to cubic complexity in classical principal component analysis. The potential of enhanced speed is suggestive that QML could outperform classical methods when analyzing high-dimensional data. In practice, the benefits remain relative since these advantages have competing realities in the overhead of mapping data, hardware noise, and limited qubits to efficiently scale the current implementations as they currently exist.

Data Efficiency

In most fields, classical machine learning needs large datasets for robust generalization, especially in deep networks where training is done with millions of parameters. A QML paradigm offers the promise of achieving meaningful performance on smaller datasets by embedding classically high-dimensional data into expressive quantum feature spaces. This becomes particularly attractive for fields such as genomics, proteomics, and medical imaging, where data is often high-dimensional, yet sample sizes are limited.

Hardware and Practicality

Classical machine learning has had decades in the spotlight to be optimized and to scale predictably on CPUs, GPUs, and distributed computing clusters. QML, meanwhile, suffers from the limitations posed by the current generation of quantum hardware, which is referred to as NISQ. These devices have a restricted number of qubits, very short coherence times, and such high error rates that large-scale practical deployments are currently infeasible.

Algorithmic Maturity

Classical machine learning algorithms, such as support vector machines, neural networks, and ensemble methods, are very well studied, and efficient implementations of these algorithms exist in standardized libraries. Compared with these, QML is a relatively young field. Considering the algorithms - but QSVMs, variational quantum circuits, and quantum principal component analysis - have shown promise, but they have not yet really outperformed strong classical baselines on large real-world data sets. Thus, QML is still an exploratory field and has limited production relevance.

Near-Term Outlook

Since classical machine learning is scalable and trustworthy, it will continue to be the primary vehicle for data analysis in real problems in the near future. Conversely, with its new applications, QML somehow stands complementary to it. The most promising applications range from hybrid workflows, where classical pre-processing and modelling are assisted by quantum subroutines for dimensionality reduction, kernel evaluation, or combinatorial optimization. So, once quantum hardware matures, QML will have great opportunities to beat classical ML in the extremely-high-dimensional problems or ones that are computationally intractable.

Implementation and reproducibility

All of the experiments were done using Python, where classical baselines were implemented with scikit-learn and quantum models with Qiskit/Qiskit Machine Learning. The PIDD was retrieved from an Excel file, and the dependent variable ("Outcome") was considered as the binary label. The zero values of the features that were physiologically implausible (glucose, blood pressure, skin thickness, insulin, and body mass index) were regarded as missing and replaced with NaN, which was followed by median imputation. For the supervised learning experiments, the data was divided into training and testing sets by a stratified 80/20 split using a fixed random seed (random_state = 42). The classical SVM model was trained with an RBF kernel on a pipeline including median imputation and standardization. The scaling of the features to (0,1) was done using MinMaxScaler for the quantum models, and then they were mapped to rotation angles in $[0, \pi]$ (angle encoding). QSVM was created with a ZZFeatureMap (reps = 2, full entanglement) and a fidelity-based statevector kernel, where training was done through QSVC. As an estimator, QNN was executed in the following way: a ZZFeatureMap (reps = 2) and RealAmplitudes ansatz (reps = 2, full entanglement) were used in combination, and optimization was done through COBYLA (maxiter = 500) with labels mapped to ±1. VQC was conducted with a statevector-based expectation-value model plus an RY encoding layer and a shallow entangling ansatz, where SPSA was the training method. QK-Means was measured through a clustering pipeline that included median imputation and standardization, and there was cluster-to-label mapping for accuracy estimation. QBM was assessed as a generative proof-of-concept through the usage of a binarized feature subset (n = 4 qubits), with SPSA optimization and evaluation by comparison metrics of distribution (e.g., total variation distance and KL divergence).

The experimental results highlight the promise and current limitations of QML applications. QSVM was able to perform with around 65% accuracy but utterly failed to classify the minority class, demonstrating a strong bias toward the majority class. The QNN, on the contrary, performed a little worse overall in terms of accuracy, at around 61%, yet it was more balanced since it partly identified the minority samples, thus showing greater flexibility at the expense of accuracy. Conversely, the VQC produced the best results, with accuracies during training and testing exceeding 71%, thus evidencing stable generalization and showing the power of the variational approach in quantum supervised learning tasks. For unsupervised learning, QK-Means got close to 69%, thus showing that quantum-enhanced similarity measures may be able to capture structures in the data that classical clustering may not. The QBM managed to reproduce a few patterns of the empirical data yet diverged considerably, indicating the difficulty of training generative QML under present noise-prone hardware conditions. This set of findings is consistent with the broader understanding where QML models, especially VQC and quantum clustering, have shown promising results but are currently limited by noise, few qubits, and difficulty in optimization.

Conclusions

QML, in comparison to classical ML, is still at an exploratory stage, but it offers distinctive promise for biomedical datasets with high dimensionality and small sample sizes, where classical methods struggle. Furthermore, conventional machine learning systems continue to be limited by scalability constraints, challenges in handling high-dimensional and imbalanced datasets, and difficulties in achieving reliable interpretability and integration into clinical workflows. These ongoing limitations reinforce the importance of exploring quantum-enhanced approaches that may offer improved representational efficiency and robustness in future medical data applications. The tendency for QML models to show imbalance across classes is not unique to the PIDD dataset but reflects broader limitations in today's hardware and algorithmic design. Current devices have very few qubits, extremely short coherence times, and high noise rates, restricting circuit depth and the amount of information preserved during encoding and measurement. Algorithmically, most quantum classifiers remain proof-of-concept, with shallow feature maps, unweighted cost functions, and simple optimizers. These constraints bias optimization in favor of majority-class solutions, leaving minority groups underrepresented. These are the broad terms in which today's QML implementations can be characterized, a situation unlikely to change until hardware scales and more robust algorithms are introduced.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Swathi Ganesan, Nalinda Somasiri

Critical review of the manuscript for important intellectual content: Swathi Ganesan, Lakmali Karunarathne, Nalinda Somasiri

Supervision: Swathi Ganesan

Acquisition, analysis, or interpretation of data: Lakmali Karunarathne, Kavindu Karunarathne, Mishani Weerakoddi

Drafting of the manuscript: Lakmali Karunarathne, Kavindu Karunarathne, Mishani Weerakoddi

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue. **Animal subjects:** All authors have confirmed that this study did not involve animal subjects or tissue. **Conflicts of interest:** In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

Data Availability Statements

Acknowledgements

The dataset used in this study is the publicly available Pima Indians Diabetes Dataset, which can be accessed via Kaggle at: https://www.kaggle.com/datasets/uciml/pima-indians-diabetes-database? The source code and analysis scripts developed for this study are stored in the YSJ RAYDAR research data repository and are available upon reasonable request from the authors. All data used in the study are de-identified.

References

1. Khaled O, Elsherif A, Salama A, Herajy M, Elsedimy E: Evaluating machine learning models for predictive analytics of liver disease detection using healthcare big data. International Journal of Electrical and Computer Engineering (IJECE). 2025, 15:1162-1174. 10.11591/ijece.v15i1.pp1162-1174
2. Esteva A, Robicquet A, Ramsundar B, et al.: A guide to deep learning in healthcare. Nature Medicine. 2019, 25:24-29. 10.1038/s41591-018-0316-z
3. Shaheen MY: Applications of Artificial Intelligence (AI) in healthcare: A review. ScienceOpen. 2021, 1-9. 10.14293/S2199-1006.1.SOR-.PPVRY8K.v1
4. LeCun Y, Bengio Y, Hinton G: Deep learning. Nature. 2015, 521:436-444. 10.1038/nature14539
5. Miotto R, Wang F, Wang S, Jiang X, Dudley JT: Deep learning for healthcare: review, opportunities and challenges. Briefings in Bioinformatics. 2017, 19:1236-1246. 10.1093/bib/bbx044
6. Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N, Lloyd S: Quantum machine learning. Nature. 2017, 549:195-202. 10.1038/nature23474
7. Dunjko V, Briegel HJ: Machine learning & artificial intelligence in the quantum domain: a review of recent progress. Reports on Progress in Physics. 2018, 81:074001. 10.1088/1361-6633/aab406
8. Schuld M, Petruccione F: Supervised Learning with Quantum Computers. Springer, Cham; 2018. 10.1007/978-3-319-96424-9
9. Bhat HA, Khanday FA, Kaushik BK, Bashir F, Shah KA: Quantum computing: fundamentals, implementations and applications. IEEE Open Journal of Nanotechnology. 2022, 3:61-77. 10.1109/ojnano.2022.3178545
10. Dejpasand MT, Sasani Ghamsari M: Research trends in quantum computers by focusing on qubits as their building blocks. Quantum Reports. 2023, 5:597-608. 10.3390/quantum5030039
11. Shafique MA, Munir A, Latif I: Quantum computing: circuits, algorithms, and applications. IEEE Access. 2024, 12:22296-22314. 10.1109/access.2024.3362955
12. Horodecki R, Horodecki P, Horodecki M, Horodecki K: Quantum entanglement. Reviews of Modern Physics. 2009, 81:865-942. 10.1103/revmodphys.81.865
13. Sanati AA, Borzoei OR: Concise insights into quantum machine learning and its practical uses. ResearchGate. 2024, 1-9. 10.13140/RG.2.2.14216.64002
14. Preskill J: Quantum computing in the NISQ era and beyond. Quantum. 2018, 2:79. 10.22331/q-2018-08-06-79

15. Schneider J, Smalley I: What is quantum computing?. (2025). Accessed: Aug 5, 2024: https://www.ibm.com/think/topics/quantum-computing.

16. D'Souza J, Chopra K, Shaikh S: Quantum perspectives: quantum computing's conquest across the spectrum of domains. Asian Journal of convergence in Technology. 2024, 9:16-23. 10.33130/ajct.2023v09i03.003

17. Mekala B, Padhy N, Penubaka KKR: Brain tumor segmentation and detection utilizing deep learning convolutional neural networks. International Journal of Computational and Experimental Science and Engineering. 2025, 11:1176-1184. 10.22399/ijcesen.1051

18. Datta S, Bernstam EV, Roberts K: A frame semantic overview of NLP-based information extraction for cancer-related EHR notes. Journal of Biomedical Informatics. 2019, 100:103301. 10.1016/j.jbi.2019.103301

19. Pandya S, Thakur A, Saxena S, et al.: A study of the recent trends of immunology: key challenges, domains, applications, datasets, and future directions. Sensors. 2021, 21:7786. 10.3390/s21237786

20. Livingston T, Muknahallipatna S: Application of variational quantum algorithms in quantum machine learning. IEEE Nanotechnology Magazine. 2025, 19:24-32. 10.1109/mnano.2025.3586032

21. Gupta R, Wood C, Engstrom T, Pole J, Shrapnel S: A systematic review of quantum machine learning for digital health. npj Digital Medicine. 2025, 8:237. 10.1038/s41746-025-01597-z

22. Ur Rasool R, Ahmad HF, Rafique W, Qayyum A, Qadir J, Anwar Z: Quantum computing for healthcare: a review. Future Internet. 2023, 15:94. 10.3390/fi15030094

23. Tijare P, Mehta NK, Raghava GPS: Benchmarking of quantum SVM and classical ML algorithms for prediction of therapeutic proteins [PREPRINT]. bioRXiv. 2025, 1:10.1101/2025.04.30.651419

24. Naresh VS, Reddi S: Quantum SVM-based predictive analytics: transforming classification methods in healthcare and beyond. Quantum Information Processing. 2025, 24:272. 10.1007/s11128-025-04892-6

25. Houssein EH, Abohashima Z, Elhoseny M, Mohamed WM: Hybrid quantum-classical convolutional neural network model for COVID-19 prediction using chest X-ray images. Journal of Computational Design and Engineering. 2022, 9:343-363. 10.1093/jcde/qwac003

26. Wang A, Mao D, Li X, Li T, Li L: HQNet: A hybrid quantum network for multi-class MRI brain classification via quantum computing. Expert Systems with Applications. 2025, 261:125537. 10.1016/j.eswa.2024.125537

27. Astuti Y, Raymant M, Quaranta V, et al.: Efferocytosis reprograms the tumor microenvironment to promote pancreatic cancer liver metastasis. Nature Cancer. 2024, 5:774-790. 10.1038/s43018-024-00731-2

28. Akpinar E, Islam SMN, Oduncuoglu M: Evaluating the impact of different quantum kernels on the classification performance of support vector machine algorithm: a medical dataset application [PREPRINT]. arXiv. 2024, 10.48550/arXiv.2407.09930

29. Ganesan S, Somasiri N: Navigating the integration of machine learning in healthcare: challenges, strategies, and ethical considerations. Journal of Computational and Cognitive Engineering. 2025, 4:8-23. 10.47852/bonviewJCCE42023600

30. Ganesan S, Somasiri N: Mitigating data scarcity in healthcare through Wasserstein generative adversarial network-a case study in medical application. 2025 10th International Conference on Machine Learning Technologies (ICMLT), Helsinki, Finland. 2025, 278-286. 10.1109/ICMLT65785.2025.11193164

31. Ganesan S, Mapa Senavige LSK, Karunarathne K, Weerakoddi M, Somasiri N: Source Code and Dataset. York St John University. Software. (2025). https://doi.org/10.25421/yorksj.30589208.v1.

32. Yin T: Quantum support vector machines: theory and applications. Theoretical and Natural Science. 2024, 51:34-42. 10.54254/2753-8818/51/2024ch0158

33. Beer K, Bondarenko D, Farrelly T, Osborne TJ, Salzmann R, Scheiermann D, Wolf R: Training deep quantum neural networks. Nature Communications. 2020, 11:808. 10.1038/s41467-020-14454-2

34. Lloyd S, Mohseni M, Rebentrost P: Quantum algorithms for supervised and unsupervised machine learning [PREPRINT]. arXiv. 2013, 10.48550/arXiv.1307.0411

35. Cerezo M, Sone A, Volkoff T, Cincio L, Coles PJ: Cost function dependent barren plateaus in shallow parametrized quantum circuits. Nature Communications. 2021, 12:1791. 10.1038/s41467-021-21728-w

36. Amin MH, Andriyash E, Rolfe J, Kulchytskyy B, Melko R: Quantum Boltzmann machine. Physical Review X. 2018, 8:1-11. 10.1103/physrevx.8.021050

How to cite this article:
Ganesan S, Karunarathne L, Karunarathne K, et al. (May 05, 2026) Quantum Artificial Intelligence for Medical Data: Comparative Evaluation of Quantum Machine Learning Models. Cureus J Comput Sci 3 : es44389-025-00063-x. DOI https://doi.org/10.7759/s44389-025-00063-x