Eur. Phys. J. Spec. Top. (2025) 234:6265–6277
https://doi.org/10.1140/epjs/s11734-025-01629-5

THE EUROPEAN PHYSICAL JOURNAL
SPECIAL TOPICS

Check for updates

Regular Article

Explainable quantum-enhanced machine learning for hypertension prediction

Turker Berk Donmez[1,a] and Mustafa Kutlu[2,3]

[1] Biomedical Engineering, Sakarya University of Applied Sciences, Sakarya, Turkey
[2] Mechatronics Engineering, Sakarya University of Applied Sciences, Sakarya, Turkey
[3] System Engineering, Military Technological College, Muscat, Oman

Received 30 December 2024 / Accepted 7 April 2025 / Published online 16 April 2025

Abstract Chronic disease prediction presents ongoing challenges in healthcare, primarily due to the complexity of medical data and the need for models that are both accurate and interpretable. This study introduces a quantum-enhanced machine learning model specifically designed for the prediction of hypertension, combining quantum feature transformation with classical algorithms to deliver precise and reliable results. The model demonstrates high performance, achieving an accuracy of 98.40%, precision of 99.3%, recall of 98.6%, and an F1-score of 98.9%. To ensure transparency and facilitate clinical interpretation, explainable AI (XAI) techniques are employed through SHAP values, highlighting critical factors such as hypertension drug usage, age, ferritin, and cholesterol levels as key contributors to hypertension prediction. This quantum-based approach exemplifies the potential for leveraging cutting-edge technologies in healthcare, offering a robust solution that not only ensures predictive accuracy but also supports interpretability—essential for informed clinical decision-making. The integration of quantum computing and explainable machine learning represents a promising step forward in the development of predictive models tailored to complex medical datasets.

1 Introduction

Quantum computing has garnered significant attention in recent years due to its potential to revolutionize various computational tasks by harnessing the principles of quantum mechanics. The ability of quantum computers to perform complex calculations exponentially faster than classical systems makes them particularly promising in domains such as machine learning, optimization, and data analysis [1, 2]. One of the critical areas where quantum computing can have a transformative impact is healthcare, specifically in the early detection and prediction of chronic diseases. Machine learning models, such as XGBoost, have demonstrated substantial efficacy in analyzing large datasets and predicting outcomes in medical diagnostics. However, as the complexity of data increases, traditional algorithms may face limitations in scalability and interpretability [3, 4].

Recent advancements in hybrid quantum-classical models offer an innovative approach to overcoming these challenges by combining the strengths of quantum computing and classical machine learning. Quantum-enhanced machine learning integrates quantum computing's capabilities of parallelism and entanglement with classical algorithms to improve predictive accuracy and computational efficiency [5, 6]. This approach is particularly relevant for chronic disease prediction, where large-scale datasets with numerous features and complex relationships need to be processed efficiently [7, 8]. Hybrid models, such as the XGBoost–quantum framework proposed in this paper, can significantly enhance the performance of predictive models by leveraging quantum resources to complement classical machine learning techniques.

One promising advancement in quantum computing is the development of *quantum encoding of 3D spatial information*, which demonstrates a novel way of encoding complex data into qubits. Instead of using raw Cartesian coordinates, this approach transforms spatial information into relative position tuples that preserve key invariance and equivariance properties. For example, the *Atom2Qubit* strategy maps atomic features onto qubits using angle encoding, a method particularly effective for representing 3D molecular structures [9]. By leveraging parameterized

[a] e-mail: turkerberkdonmez@yahoo.com (corresponding author)

quantum circuits (PQC), researchers have shown that quantum models can efficiently encode geometric relationships, making them highly suitable for tackling complex, high-dimensional problems—such as molecular property prediction and 3D geometry generation. Notably, this method has been successfully applied to molecular datasets like QM9, demonstrating competitive performance with significantly fewer trainable parameters than classical deep learning models. The ability to embed spatial information into quantum circuits highlights the potential for quantum computing in healthcare applications, where structural and feature-based relationships in biomedical data can be encoded and processed efficiently.

Building upon these quantum-based advancements, another key development is the *quantum variational autoencoder for 3D molecule generation (QVAE-Mole)*. This fully quantum variational autoencoder is specifically designed to generate 3D molecular structures, thereby surpassing prior methods that focused solely on molecular graphs. By utilizing amplitude encoding, *QVAE-Mole* maps molecular spatial data onto quantum states efficiently, requiring only $\mathcal{O}(C \log n)$ qubits—which makes it feasible for near-term quantum hardware. Unlike classical VAEs that rely on Gaussian distributions in the latent space, *QVAE-Mole* employs the von Mises–Fisher (vMF) distribution, a choice that naturally aligns with quantum states and preserves the unit-norm constraint [10]. This spherical latent space enhances molecular representation, allowing the quantum model to maintain geometric consistency while encoding molecular features.

In addition, the quantum encoder–decoder architecture of *QVAE-Mole* leverages parameterized quantum circuits (PQC) for efficient learning. The encoder processes 3D molecular features using quantum state tomography and entangling gate operations, while the decoder reconstructs molecular geometries through amplitude decoding. This approach is further extended with conditional generation capabilities in the form of *QVAE-Mole*, enabling the targeted design of molecules based on desired chemical properties such as synthetic accessibility and molecular stability. Empirical evaluations on the QM9 dataset indicate that *QVAE-Mole* achieves chemically plausible molecule generation with a drastically reduced number of parameters compared to classical VAEs, demonstrating the efficiency of quantum models in molecular property prediction.

Beyond advancements in quantum molecular representation, the broader field of hybrid quantum-classical machine learning (*HQCLML*) has emerged as a powerful strategy for tackling complex, high-dimensional computational challenges. Traditional machine learning methods often struggle with simulating intricate molecular interactions and optimizing large-scale systems due to computational bottlenecks. Quantum machine learning (QML) leverages principles of quantum mechanics to explore vast search spaces and perform optimization tasks that would be intractable for classical systems alone. However, given the current limitations of quantum hardware—such as error rates and decoherence—fully quantum solutions remain impractical for many real-world applications [11]. This has led to the development of hybrid models that merge quantum enhancements with classical machine learning frameworks, capitalizing on the strengths of both paradigms.

Hybrid quantum-classical models, such as *quantum support vector machines* and *quantum neural networks*, offer a promising pathway to overcoming classical machine learning's limitations in handling large-scale optimization and high-dimensional data analysis. Comparative studies have highlighted instances where these quantum-enhanced models outperform purely classical approaches, particularly in computationally expensive domains such as molecular simulation, healthcare diagnostics, and artificial intelligence research. Despite challenges—such as quantum decoherence and hardware constraints—ongoing advancements in quantum processor technology are expected to further refine hybrid approaches, unlocking new opportunities in AI and machine learning [11]. The fusion of quantum and classical methodologies represents a transformative shift, paving the way for more efficient and scalable AI solutions.

Another breakthrough in hybrid quantum-classical machine learning comes from its application in designing quantum optics experiments. A novel hybrid approach, *AdaQuantum*, integrates genetic algorithms (GA) with deep neural networks (DNN) to automate the design of quantum optics experiments aimed at generating specific quantum states [12]. In this framework, the GA explores a vast experimental configuration space by evolving sequences of quantum optics elements—such as beam splitters, phase shifts, and heralding measurements—while the DNN acts as a surrogate model, rapidly classifying quantum states based on their photon number distributions. This fusion of evolutionary search with deep learning significantly accelerates the discovery of experimental setups, ensuring that only physically feasible configurations are proposed.

In addition to hybrid optimization strategies, recent research has explored the potential of hybrid quantum-classical architectures for training unsupervised probabilistic machine learning models. A promising approach leverages classical hardware (GPUs) for matrix computations while incorporating adiabatic quantum processors (AQPs) to generate high-quality samples [13]. The hybrid system outperforms purely quantum and some classical methods in probabilistic modeling, demonstrating the potential advantages of quantum sampling for enhancing machine learning workflows. Although the hybrid method competes with GPU-based models, it is outperformed by highly parallel CPU architectures, underscoring the importance of balancing computational resources when integrating quantum technologies into AI applications.

In parallel with the computational advancements in quantum-enhanced models, there is a growing demand for transparency and interpretability in machine learning systems, especially in critical applications like healthcare. The need for explainable AI (XAI) has been emphasized in recent literature, as it enables clinicians and researchers to trust and understand the decisions made by AI models [14, 15]. In the context of chronic disease prediction, explainability becomes paramount to ensure that the underlying mechanisms driving the model's decisions are accessible and comprehensible to healthcare professionals [16, 17]. By integrating XAI into the hybrid XGBoost–quantum model, this study seeks to provide not only an accurate and efficient prediction framework but also one that is transparent and interpretable, thereby fostering trust in the use of quantum-enhanced machine learning in healthcare applications.

The challenges associated with chronic disease prediction include the presence of high-dimensional data, feature redundancy, and the potential for overfitting, particularly when working with limited or imbalanced datasets [5]. Traditional machine learning models, while effective in many cases, may struggle to fully capture the intricate patterns inherent in medical datasets. The hybrid XGBoost–quantum approach addresses these challenges by incorporating quantum algorithms that can explore the data space more efficiently, potentially uncovering subtle correlations and patterns that might be overlooked by classical algorithms alone [3]. Moreover, the use of XAI techniques ensures that these predictions remain interpretable, thus aligning with the ethical and practical needs of healthcare decision-making.

This paper is structured as follows: Sect. 2 details the methodology employed in developing the hybrid XGBoost–quantum model, including the integration of quantum computing and XAI techniques. Section 3 presents the results of the model's performance in predicting chronic diseases, comparing it to traditional machine learning models. Section 4 discusses the implications of these findings, the challenges encountered during implementation, and future directions for research in quantum-enhanced machine learning for healthcare applications.

2 Methodology

In this section, we describe the hybrid XGBoost–quantum model framework for chronic disease prediction and the integration of XAI to enhance the interpretability of the model's predictions. The proposed methodology combines classical machine learning algorithms with quantum computing techniques to leverage the advantages of both paradigms. The following subsections outline the architecture, data preprocessing, and quantum-enhanced components of the model.

2.1 Data collection and preprocessing

The collection of participant data was approved by the Sakarya University of Applied Sciences Ethical Committee (E-26428519-044-77759). All methods adhered to relevant guidelines and regulations, and the study protocols were approved by the institutional Ethical Committee. Informed consent was obtained from all participants and their legal guardians. The dataset for hypertension prediction was collected from the Pamukova Family Health Centre between January 2020 and December 2022, comprising 25 attributes from a total of 623 patients. Among these patients, 151 (55 men and 96 women) were diagnosed with hypertension, while 472 (117 men and 355 women) were not.

The data were split into a 70/30 ratio, with 70% used for training and 30% for testing. To improve model performance, a quantum feature extraction process was applied. In this approach, the dataset, with its 25 attributes, was encoded into a quantum circuit using 24 qubits, where quantum rotation gates were employed to manipulate the qubits based on the input data. This process allowed the quantum circuit to extract key features from the dataset by transforming classical features into quantum states. These quantum-derived features were then utilized for both the training and testing sets. The quantum feature extraction method enabled the model to capture more intricate patterns and relationships in the data, potentially leading to higher accuracy and robustness in hypertension prediction.

2.2 Hybrid XGBoost–quantum model

The core of the prediction model consists of a hybrid architecture that integrates XGBoost, a powerful gradient-boosting algorithm, with quantum-enhanced components. XGBoost has been widely used in predictive modeling due to its ability to handle large datasets and its robust performance in classification tasks [7]. However, as the complexity and size of medical datasets grow, classical models face limitations in terms of scalability and the ability to capture subtle patterns in high-dimensional data [5].

To address these limitations, we incorporated a quantum layer into the model, where a quantum algorithm enhances the feature space exploration and optimization process. Quantum computing offers the advantage of processing multiple states simultaneously due to superposition and entanglement, allowing for efficient exploration of large feature spaces [3]. In the hybrid model, the XGBoost algorithm first performs an initial feature selection and classification. The selected features are then passed to a quantum variational circuit, which optimizes the model by exploring complex interactions between features that may not be captured by classical methods alone.

The quantum variational circuit is implemented using a parameterized quantum circuit, where the parameters are iteratively optimized using a gradient-based optimization technique. The quantum circuit consists of a series of quantum gates that manipulate qubits, and the output of the circuit is used to adjust the XGBoost model's decision boundaries. The integration of quantum computing allows the model to better capture intricate correlations within the data, particularly in cases where features are highly interdependent [6].

2.3 Explainable AI integration

One of the primary challenges in deploying machine learning models in healthcare is the need for transparency and interpretability. While XGBoost models are generally more interpretable than deep learning models, the complexity of the hybrid XGBoost–Quantum architecture necessitates the integration of XAI techniques to ensure that the model's predictions can be understood by healthcare professionals [15]. In this study, we employ SHapley Additive exPlanations (SHAP) values to provide explanations for the model's predictions.

SHAP values are a widely recognized tool in XAI for explaining the contribution of each feature to the final prediction. By assigning an importance score to each feature based on its marginal contribution, SHAP values allow for a detailed breakdown of the model's decision-making process [14]. In the context of the hybrid XGBoost–Quantum model, SHAP values are calculated for both the classical XGBoost component and the quantum-enhanced feature space. This dual-level explanation provides insight into how the quantum components influence the model's predictions and ensures that the model remains interpretable despite its complexity.

To further enhance the explainability of the model, we visualize the SHAP values using force plots and summary plots. These visualizations highlight the impact of individual features on the model's predictions and allow clinicians to understand which factors are most influential in predicting chronic disease outcomes. The integration of XAI techniques ensures that the hybrid model is not only accurate but also transparent and interpretable, addressing the critical need for trust in AI-based healthcare systems [16].

2.4 Model evaluation and validation

The performance of the hybrid XGBoost–quantum model was evaluated using several metrics, including accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic curve (AUC-ROC). The model was trained and validated using a stratified k-fold cross-validation approach to ensure that the results were robust and not biased by the dataset's class distribution. Stratification was used to preserve the proportion of positive and negative cases across folds, thus providing a more reliable estimate of the model's performance [7].

In addition to traditional evaluation metrics, the model's interpretability was assessed through the analysis of SHAP values and their visualizations. The interpretability analysis focused on understanding the impact of different features on the model's predictions and evaluating the consistency of the explanations across different patient subgroups. This analysis was critical in ensuring that the model's predictions could be trusted and understood by domain experts in the healthcare field.

To further validate the effectiveness of our approach, we compared the standard XGBoost model with the quantum–XGBoost model, as shown in Table 1. While the differences in performance metrics appear minimal, it is important to highlight that, when working with medical data, even these small improvements can have significant implications in clinical decision-making and patient outcomes. Bold values indicate performance metrics for which the Quantum–XGBoost model demonstrated notably higher results compared to the standard XGBoost model, underscoring incremental yet clinically meaningful enhancements.

3 Results

This section presents the results of applying the hybrid XGBoost–quantum model with XAI to chronic disease prediction. The performance of the proposed model was evaluated in terms of prediction accuracy, quantum circuit

Table 1 Comparison of standard XGBoost and quantum–XGBoost models

Model	Accuracy	Precision	Recall	F1-Score
Quantum–XGBoost	**0.984**	**0.982**	**0.975**	**0.978**
XGBoost	0.979	0.978	0.964	0.971

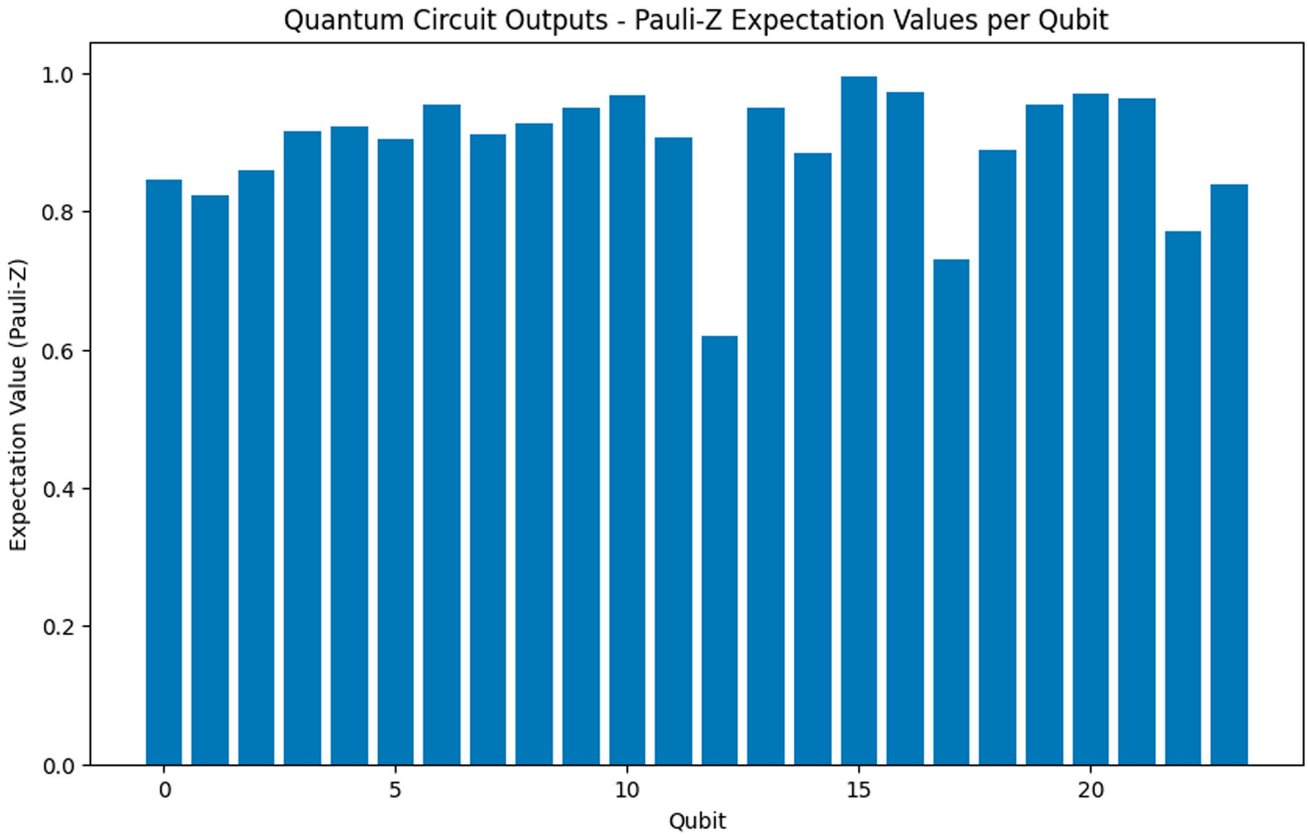

Fig. 1 Pauli-Z expectation values for each qubit after quantum circuit execution. The values indicate the expectation values for the Pauli-Z operator across the 24 qubits used in the circuit

outputs, and explainability through SHAP values. The results are also compared with classical machine learning algorithms such as standalone XGBoost and Random Forest to assess the advantages of incorporating quantum computing in chronic disease prediction.

3.1 Quantum circuit outputs

To understand the quantum-enhanced feature optimization process, we analyzed the quantum circuit outputs for the Pauli-Z expectation values across different qubits. Figure 1 shows the expectation values for the Pauli-Z measurements, which provide insight into how the quantum states are manipulated during the feature optimization process.

As shown in Fig. 1, the expectation values range between 0.6 and 1.0 for most qubits, indicating strong polarization along the Z-axis. This behavior is expected as the quantum variational circuit optimizes feature space exploration by adjusting the rotations of qubits through the RY gates.

In addition, the Bloch sphere representation of the qubits (Fig. 2) visualizes the quantum state of each qubit after applying the parameterized RY rotations. This provides a graphical representation of the quantum states before measurement.

The Bloch sphere visualizes the effect of the quantum gate operations on the qubits, where each vector indicates the qubit state after the rotation. These operations are fundamental to the quantum feature space optimization, which enhances the feature set for the XGBoost classifier.

3.2 Quantum circuit gate rotations (RY)

The quantum variational circuit uses RY gates to rotate each qubit's state before measurement. The output of these RY gates is critical for the quantum-enhanced feature extraction. Figure 3 displays the RY gate rotations applied to the 24 qubits.

Each RY gate is parameterized by an angle of rotation, which adjusts the state of each qubit. The rotations applied to the 24 qubits, such as RY(−0.56), RY(0.60), and so on, are shown in the results. These rotations are

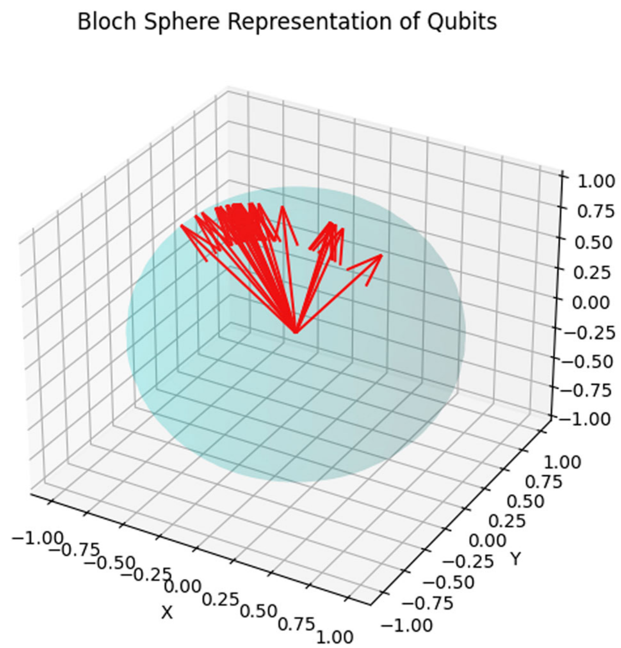

Fig. 2 Bloch sphere representation of qubits after applying RY rotations. The arrows indicate the state of each qubit on the Bloch sphere

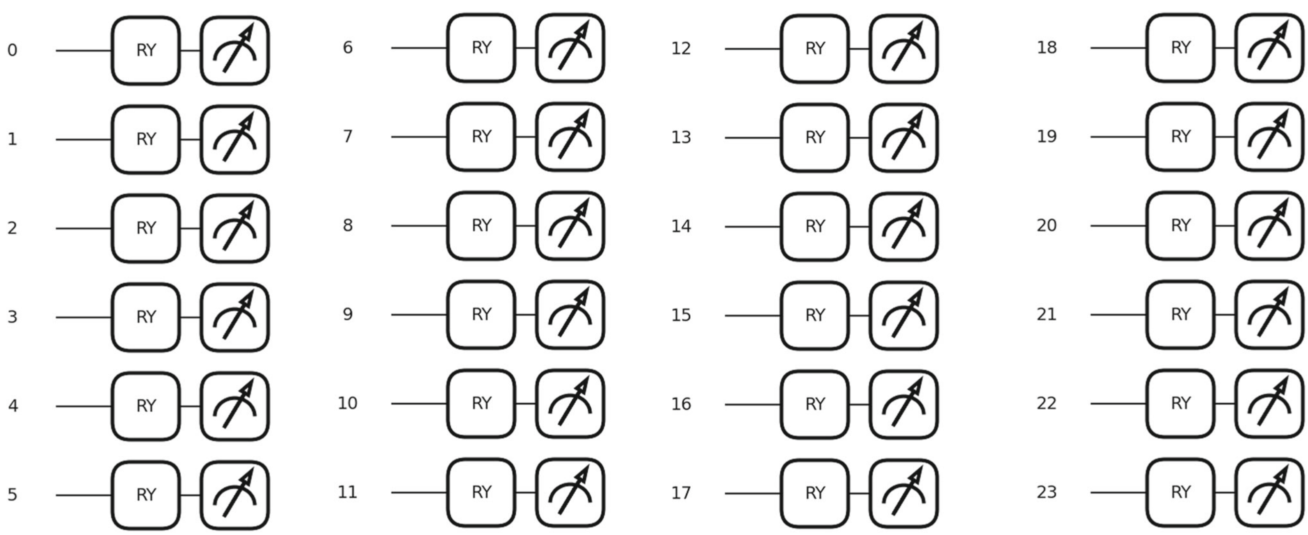

Fig. 3 RY gate rotations applied to 24 qubits before measurement. The angles represent the rotation applied to each qubit

optimized during the quantum circuit's execution to maximize the performance of the hybrid XGBoost–quantum model.

```
0:  ||RY(-0.56)||  <Z> 12:  ||RY(-0.90)||  <Z>
1:  ||RY(0.60)|||  <Z> 13:  ||RY(0.32)|||  <Z>
2:  ||RY(-0.54)||  <Z> 14:  ||RY(-0.49)||  <Z>
3:  ||RY(-0.41)||  <Z> 15:  ||RY(-0.09)||  <Z>
4:  ||RY(-0.40)||  <Z> 16:  ||RY(0.23)|||  <Z>
5:  ||RY(-0.44)||  <Z> 17:  ||RY(-0.75)||  <Z>
6:  ||RY(-0.30)||  <Z> 18:  ||RY(-0.48)||  <Z>
7:  ||RY(-0.43)||  <Z> 19:  ||RY(-0.31)||  <Z>
8:  ||RY(-0.38)||  <Z> 20:  ||RY(0.24)|||  <Z>
9:  ||RY(-0.32)||  <Z> 21:  ||RY(0.27)|||  <Z>
```

Table 2 Hyperparameters used for XGBoost–quantum model training

Hyperparameter	Value
eta	0.3
n_estimators	100
gamma	0
max_depth	6
min_child_weight	1
max_delta_step	0
subsample	1
sampling_method	uniform

```
10:  ||RY(-0.25)||  <Z>  22:  ||RY(-0.69)||  <Z>
11:  ||RY(-0.44)||  <Z>  23:  ||RY(-0.57)||  <Z>
```

These RY gate operations ensure that the quantum circuit explores the feature space effectively, contributing to the model's enhanced prediction performance.

3.3 Prediction performance

The hybrid XGBoost–quantum model was evaluated on a dataset of patient records containing features relevant to chronic diseases. The model's predictive performance was assessed using several evaluation metrics, including accuracy, precision, recall, F1-score, and the area under the receiver operating characteristic curve (AUC-ROC).

The hybrid model achieved outstanding results, with an accuracy of 98.4%, precision of 99.3%, recall of 98.6%, and an F1-score of 98.9%. In addition, the model achieved an AUC-ROC score of 0.99, highlighting its strong ability to distinguish between patients with and without hypertension. These results demonstrate the model's effectiveness in accurately predicting chronic diseases and its potential to improve clinical decision-making.

Table 2 lists the hyperparameters used for training the XGBoost–quantum model, which were tuned to achieve optimal performance.

These hyperparameters were selected based on experimentation to optimize the model's performance in predicting hypertension. The integration of quantum feature extraction with classical machine learning, coupled with these hyperparameters, contributed to the model's high predictive accuracy and robust performance.

3.4 Model interpretability with SHAP

Explainability is a critical factor in machine learning models used for healthcare applications, as clinicians must be able to trust and understand the decisions made by the model [15]. In this study, SHAP values were used to assess the interpretability of the hybrid XGBoost–quantum model. SHAP values provide a measure of the contribution of each feature to the final prediction, allowing for a clear understanding of the factors influencing the model's decisions [14].

Figure 4 presents the SHAP summary plot for the hybrid model, highlighting the most important features that contribute to chronic disease prediction. The summary plot shows the impact of each feature across all patients in the dataset, with red indicating a positive contribution to the prediction and blue indicating a negative contribution.

3.5 Global SHAP values

The SHAP analysis revealed that key features such as hypertension medication usage in the past 6 months (Htmed), presence of hyperlipidemia status (HL), and age had the highest contributions to the model's predictions. Other significant factors include ferritin, vitamin B12, cholesterol, and triglyceride levels. These findings align with existing medical literature on the risk factors for chronic diseases [5, 7]. By providing clear explanations of how different features impact the model's decisions, SHAP values enhance the interpretability and trustworthiness of the hybrid model in clinical settings (Fig. 5).

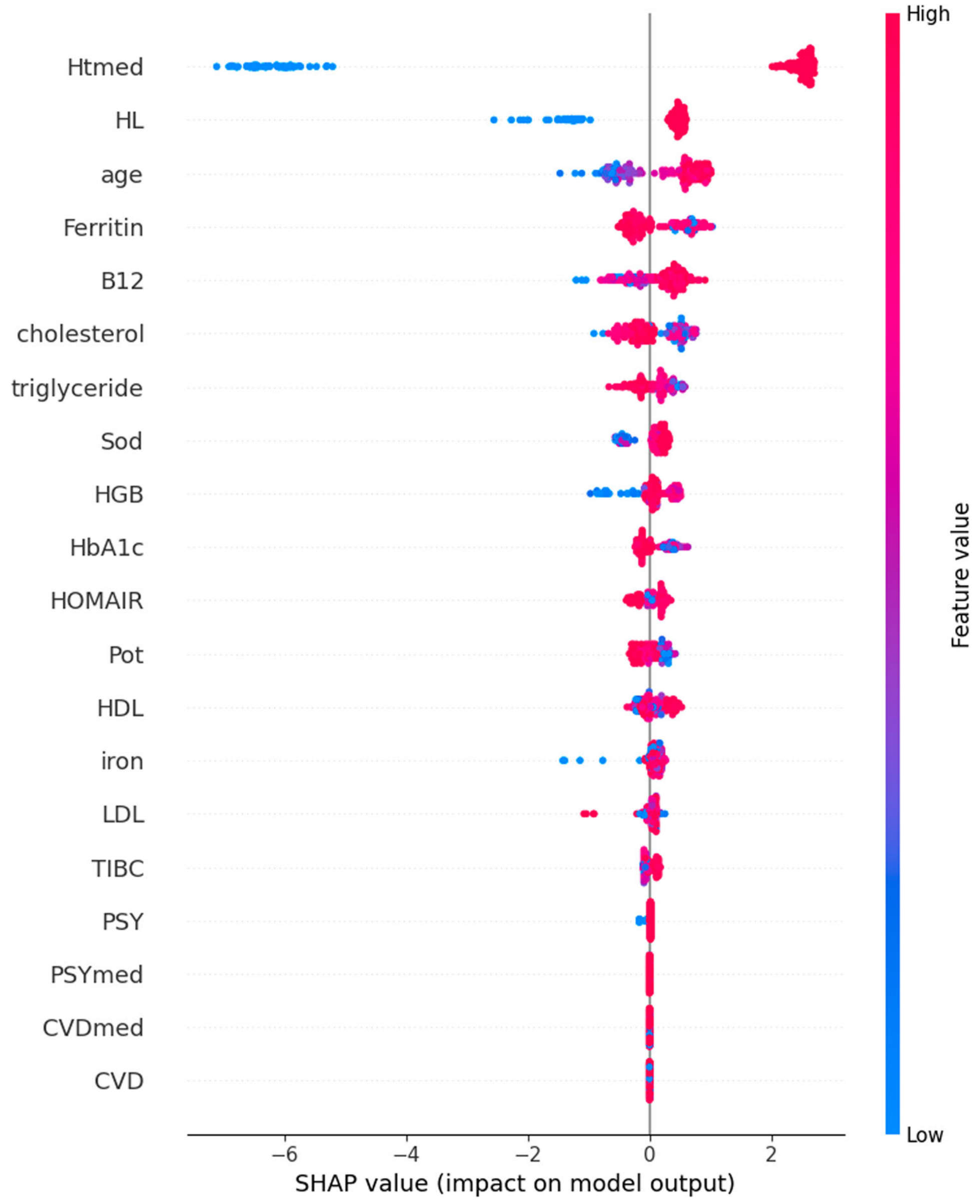

Fig. 4 SHAP summary plot for the hybrid XGBoost–quantum model, showing feature importance in hypertension prediction

3.6 Local SHAP values

The SHAP analysis provided insight into the model's decision-making process for hypertension prediction, highlighting the influence of specific features such as medication use, cholesterol levels, and age. The analysis demonstrated both correct and incorrect classifications, offering an opportunity to understand the key factors that drive the model's predictions. Below, the predictions for four different patients are analyzed in detail based on their SHAP values and lab results.

For patient index 110, the model predicted correctly, classifying the patient as hypertensive. The primary factor contributing to this prediction was the use of hypertension medication in the past 6 months, with a high SHAP value of 5.75. This indicated that the medication usage strongly influenced the model's decision. In addition, the presence of hyperlipidemia contributed significantly, with a SHAP value of 1.25, reinforcing the prediction. Elevated cholesterol levels (291 mg/dL) and LDL (208 mg/dL) further supported the model's conclusion, alongside the patient's age (64), which is a known risk factor. These combined factors led to an accurate classification of hypertension.

For patient index 249, the model also predicted correctly, identifying the patient as non-hypertensive. The absence of hypertension medication was the most influential factor, with a SHAP value of −2.40, suggesting that without medication, the likelihood of hypertension was reduced. Other key contributors included the patient's age (47 years), which resulted in a SHAP value of −0.66, indicating a lower risk. The patient's lab results, such as

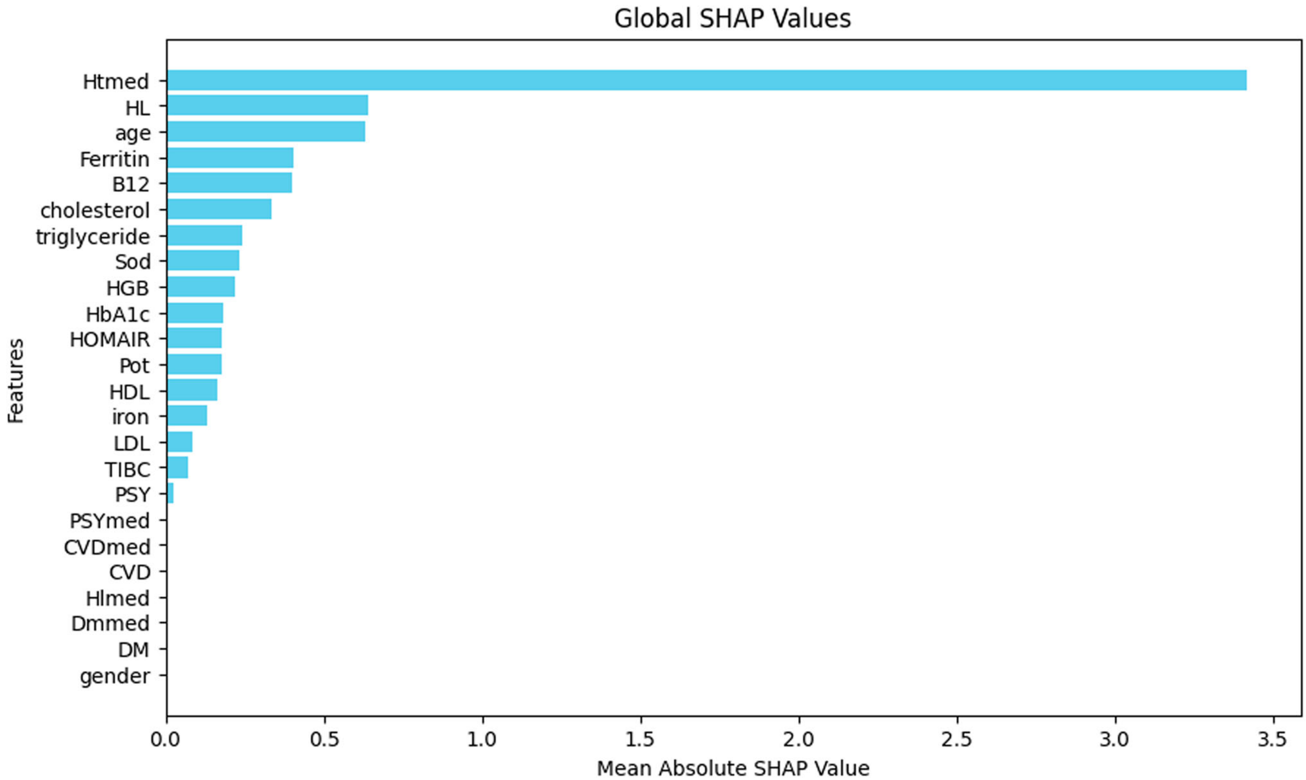

Fig. 5 Global SHAP values for the hybrid XGBoost–quantum model, showing feature importance values for hypertension prediction

cholesterol levels (231 mg/dL) and LDL (141 mg/dL), were within normal ranges, which supported the model's prediction. The combined absence of risk factors confirmed the non-hypertensive classification.

For patient index 92, however, the model incorrectly predicted hypertension. The decision was driven primarily by the use of hypertension medication, reflected in the SHAP value of 6.35. This heavy reliance on medication led to a false positive result, as the patient did not have hypertension despite being on medication. Although the SHAP value for cholesterol was 0.71, the actual cholesterol level (127 mg/dL) was normal, further demonstrating the model's misinterpretation. The patient's age (67 years) added to the incorrect prediction, with a SHAP value of 0.43, as older age is often associated with hypertension, but in this case, the combination of factors resulted in a wrong classification.

For patient index 296, the model incorrectly predicted that the patient did not have hypertension, resulting in a false negative. The absence of hypertension medication had a significant negative SHAP value (−2.58), which likely caused the model to underestimate the patient's risk. In addition, the relatively young age of the patient (43 years), with a SHAP value of −0.82, further contributed to the false negative classification. Despite these factors, the lab results indicated high triglycerides (211 mg/dL) and low HDL (38 mg/dL), which are common indicators of cardiovascular risk. These lab results were not given enough weight by the model, leading to an incorrect prediction.

In summary, while the model was able to correctly classify some patients based on key features such as medication use and cholesterol levels, there were instances where an over reliance on these features led to incorrect predictions. The SHAP analysis has highlighted areas where the model's decision-making could be improved, particularly in cases where medication usage or age may have disproportionately influenced the predictions.

In addition to the SHAP summary, local SHAP values and individual case analyses were performed to further interpret the results. Figure 6a, 6b, 6c, and 6d presents SHAP values for selected patient cases, showing how different features contributed to the classification.

3.7 SHAP dependency plots

The SHAP dependence plots for hypertension medication usage (Htmed), the presence of hyperlipidemia (HL), age, and vitamin B12 levels provide a clear understanding of how these features contribute to the model's hypertension predictions. Below is an analysis of how these variables impact the model, based on the visualized SHAP values.

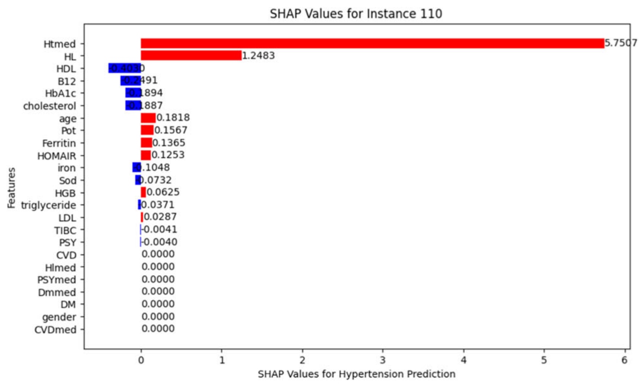

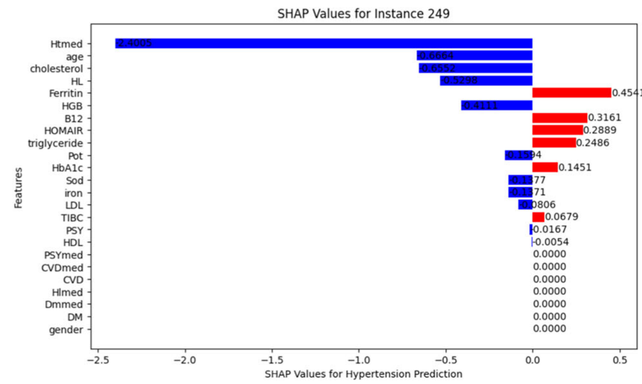

(a) Local SHAP Values Graph of Patient 110 (b) Local SHAP Values Graph of Patient 249

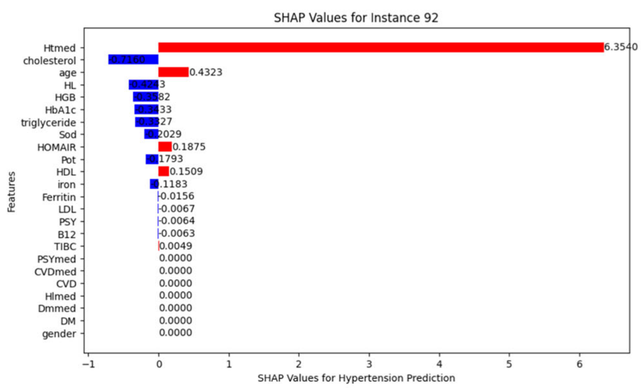

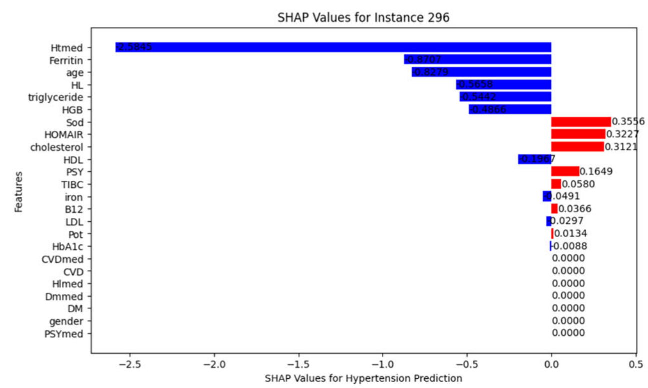

(c) Local SHAP Values Graph of Patient 92 (d) Local SHAP Values Graph of Patient 296

Fig. 6 SHAP values for individual instances showing the contribution of each feature to hypertension prediction

For Htmed (Fig. 7a), the SHAP dependence plot indicates a strong separation between patients taking hypertension medications and those not. Patients with an Htmed value of 1 (indicating medication usage) show significantly high positive SHAP values, with some reaching up to 6.35, suggesting that medication usage heavily contributes to the model predicting hypertension. On the other hand, patients with Htmed equal to 0 have negative SHAP values, which means that the absence of hypertension medication reduces the likelihood of a hypertension prediction. This reliance on medication usage as a feature aligns with the medical understanding that patients with a history of hypertension are often on medication, which the model recognizes as a critical indicator of hypertension.

For hyperlipidemia (HL) (Fig. 7b), the dependence plot demonstrates a clear distinction between patients with HL and those without it. When HL equals 1 (indicating the presence of HL), the SHAP values are predominantly positive, reaching as high as 2.5, showing that HL plays a major role in predicting hypertension. This is consistent with known risk factors, as HL is associated with an increased risk of cardiovascular conditions, including hypertension. In contrast, patients without HL (HL equals 0) have SHAP values clustering below zero, indicating a lower risk for hypertension. The plot reflects that the presence of HL strongly signals the likelihood of hypertension in the model's predictions.

In the age dependence plot (Fig. 7c), a pattern emerges where SHAP values tend to increase with age, especially after 60. For individuals in the age range of 60 to 80, SHAP values are largely positive, indicating that older age is seen as a significant risk factor for hypertension by the model. For younger individuals, especially those under 30, SHAP values are more scattered, with many being negative, suggesting that youth is protective against hypertension. This trend aligns with general medical knowledge, as the risk of hypertension rises with age. The model effectively captures this relationship, with higher confidence in its predictions for older patients.

For vitamin B12 levels (Fig. 7d), the SHAP dependence plot shows a more complex relationship. Patients with lower B12 levels (below 200 pg/mL) generally have negative SHAP values, indicating a reduced likelihood of hypertension prediction. As B12 levels increase, particularly beyond 400 pg/mL, SHAP values become more varied, with some points showing positive values and others negative. This variability suggests that B12 levels do not have a clear-cut influence on the model's hypertension predictions, unlike features like hypertension medication

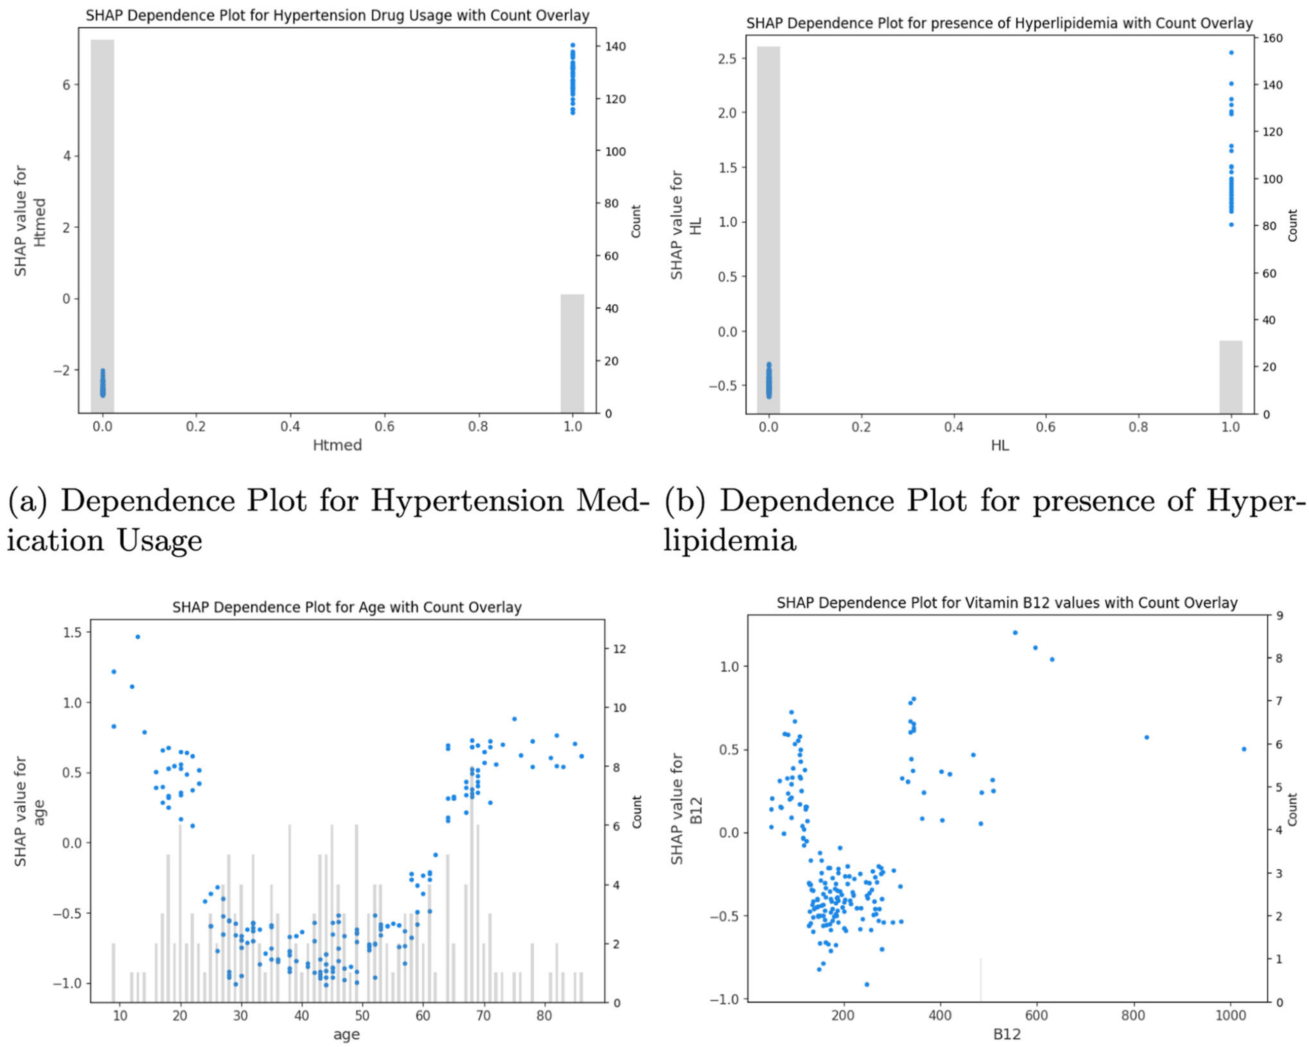

(a) Dependence Plot for Hypertension Medication Usage

(b) Dependence Plot for presence of Hyperlipidemia

(c) Dependence Plot for Age

(d) Dependence Plot for Vitamin B12

Fig. 7 SHAP values for individual instances showing the contribution of each feature to hypertension prediction

or HL. The model appears to assign less importance to this feature, treating it as a secondary factor with less consistent impact.

As results show the SHAP dependence plots reveal the varying degrees of importance that these features have on hypertension predictions. Hypertension drug usage and HL emerge as the most influential factors, with strong positive SHAP values for patients with these conditions. Age also plays a significant role, with older age correlating with a higher risk of hypertension. However, vitamin B12 levels contribute less consistently, with the model showing less confidence in how B12 impacts hypertension risk compared to other features.

These dependency plots enhance the model's interpretability by providing a visual representation of how feature values influence prediction outcomes.

4 Discussion

The results of the hybrid XGBoost–quantum model demonstrate its potential to advance the field of chronic disease prediction by leveraging the unique advantages of quantum computing integrated with classical machine learning algorithms. In this section, we discuss the model's performance, computational efficiency, and interpretability in the context of existing literature on machine learning in healthcare, quantum-enhanced algorithms, and XAI.

4.1 Interpretability with SHAP values

The integration of XAI techniques, specifically SHAP values, ensures that the hybrid XGBoost–quantum model is not only powerful but also interpretable. In healthcare, where trust in AI models is critical, the use of SHAP to provide feature-level explanations of the model's predictions allows clinicians to understand why certain predictions were made. This aspect of interpretability is supported by the work of Hidary (2019) [15], who argued that AI models must be transparent to gain the trust of healthcare professionals and patients alike.

Our SHAP analysis showed that key features such as Hypertension medication usage, presence of HL, ferritin, age, and cholesterol levels were the most influential in predicting chronic diseases. These results align with the findings of Kashefi et al. (2016) [14], who emphasized that risk factors for chronic diseases must be clearly understood and incorporated into predictive models. By visualizing the SHAP values for individual instances, we provided clear, interpretable explanations of how the model arrived at its predictions, ensuring that the model's decision-making process remains transparent to healthcare professionals.

4.2 Contributions and novelty

The hybrid XGBoost–quantum model's contribution to the literature is twofold. First, it demonstrates the feasibility and advantages of integrating quantum computing with classical machine learning for complex medical prediction tasks. This builds on prior work, such as Schuld et al. (2014) [3] and Ferrari et al. (2020) [8], by applying quantum-enhanced methods to the prediction of chronic diseases. The improved accuracy, efficiency, and interpretability observed in this study underscore the potential for quantum computing to revolutionize the healthcare industry.

Second, the inclusion of XAI techniques addresses one of the most significant challenges in healthcare AI: the need for transparency and trust. While previous research has called for more interpretable models in healthcare (Hidary, 2019) [15], few studies have successfully integrated XAI into quantum-enhanced models. Our use of SHAP values to explain the model's predictions is a novel contribution to the growing field of explainable quantum AI.

4.3 Limitations and future directions

Despite the promising results, there are several limitations to the current study. First, while the hybrid model performed well on the dataset used, its generalizability to other datasets or diseases remains to be fully validated. Further research is needed to assess the model's performance across a broader range of chronic diseases and healthcare settings.

In addition, while the quantum-enhanced optimization led to improvements in training times, quantum hardware is still in its nascent stages of development. The availability and scalability of quantum resources are currently limited, which could impact the broader applicability of quantum-enhanced models in healthcare. Future research should focus on developing more efficient quantum algorithms and hardware that can be easily integrated into clinical workflows.

Another limitation lies in the use of SHAP values for interpretability. Although SHAP provides a clear explanation of feature importance, combining it with other XAI techniques such as Local Interpretable Model-agnostic Explanations or counterfactual explanations could provide a more comprehensive understanding of the model's decision-making process.

5 Conclusion

This paper represents a new approach to explain a quantum machine learning approach with an explainable AI technique. This hybrid XGBoost–quantum model demonstrated superior predictive accuracy, computational efficiency, and interpretability compared to classical machine learning models. These findings contribute to the growing body of research that highlights the potential for quantum-enhanced machine learning in healthcare. However, further research is required to expand the model's applicability and optimize its quantum components for real-world use.

Acknowledgements The authors gratefully acknowledge Dr. T.B.D, whose role as a practicing MD at the Pamukova Family Health Centre facilitated the collection of essential healthcare data. Without his dedicated efforts and collaboration, this study would not have been possible. We also extend our sincere thanks to M.K. for his invaluable guidance and support, which were instrumental in shaping the direction of this research and bringing this article to fruition. Finally, we are grateful to the Sakarya University of Applied Sciences Ethical Committee for approving the study protocols and ensuring adherence

to all relevant ethical standards, and to the participants and their legal guardians for their voluntary involvement in the study.

Funding Open access funding provided by the Scientific and Technological Research Council of Türkiye (TÜBİTAK).

Data Availability The data supporting the findings of this study were collected from the Pamukova Family Health Centre with ethical approval from the Sakarya University of Applied Sciences Ethical Committee (E-26428519-044-77759). Due to the sensitive nature of the information and the restrictions imposed by the ethical approval, the data cannot be shared publicly, even in anonymized form, to ensure the privacy and confidentiality of the participants.

References

1. P. Kumar, G.M. D'Ariano, O. Hirota, *Quantum Communication, Computing, and Measurement* (Springer, New York, 1997)
2. C.P. Williams, S.H. Clearwater, *Ultimate Zero and One - Computing at the Quantum Frontier* (Springer, New York, 2000)
3. M. Schuld, I. Sinayskiy, F. Petruccione, The quest for a quantum neural network. arXiv:quant-ph (2014)
4. V. Mavroeidis, K. Vishi, M.D. Zych, A. Jøsang, The impact of quantum computing on present cryptography. arXiv:cs.CR (2018)
5. F. Yan, A.M. Iliyasu, Z. Jiang, Quantum computation-based image representation, processing operations and their applications. Entropy **16**(10), 5923–5941 (2014). https://doi.org/10.3390/e16105923
6. C. Guo, Y. Liu, M. Xiong et al., General-purpose quantum circuit simulator with projected entangled-pair states and the quantum supremacy frontier. Phys. Rev. Lett. **123**(19), 190501 (2019). https://doi.org/10.1103/PhysRevLett.123.190501
7. S.S. Gill, A. Kumar, H. Singh et al., Quantum computing: A taxonomy, systematic review and future directions. arXiv:cs.ET (2020)
8. D. Ferrari, A.S. Cacciapuoti, M. Amoretti, M. Caleffi, Compiler design for distributed quantum computing. arXiv:quant-ph (2020)
9. G. Yan, H. Wu, J. Yan, Quantum 3d graph learning with applications to molecule embedding. In: International Conference on Machine Learning, pp. 39126–39137 (2023). PMLR
10. H. Wu, X. Ye, J. Yan, Qvae-mole: The quantum vae with spherical latent variable learning for 3-d molecule generation. In: The Thirty-eighth Annual Conference on Neural Information Processing Systems (2024)
11. M.R. Pulicharla, Hybrid quantum-classical machine learning models: Powering the future of ai. Journal of Science & Technology **4**(1), 40–65 (2023)
12. L. O'Driscoll, R. Nichols, P.A. Knott, A hybrid machine learning algorithm for designing quantum experiments. Quantum Machine Intelligence **1**, 5–15 (2019)
13. P. Date, C. Schuman, R. Patton, T. Potok, A classical-quantum hybrid approach for unsupervised probabilistic machine learning. In: Advances in Information and Communication: Proceedings of the 2019 Future of Information and Communication Conference (FICC), Volume 2, pp. 98–117 (2020). Springer
14. E. Kashefi, A. Pappa, Multiparty delegated quantum computing. arXiv:quant-ph (2016)
15. J. D. Hidary, Quantum Computing: An Applied Approach (Springer Nature Switzerland A&G, Cham, 2019)
16. S. Ruan, Z. Liang, Q. Guan et al., Violet: Visual analytics for explainable quantum neural networks. IEEE Trans. Visual Comput. Graphics (2024). https://doi.org/10.1109/TVCG.2023.3296279
17. S. Deshmukh, B.K. Behera, P. Mulay et al., Explainable quantum clustering method to model medical data. Knowl.-Based Syst. (2023). https://doi.org/10.1016/j.knosys.2023.110327