Article

Not peer-reviewed version

Quantum-Enhanced Oral Disease Detection Using Hybrid Quantum-Classical Neural Networks

Md. Shakhawat Hossain , Md. Mehedi Hasan , Mohammad Junayed Hasan * , M. R. C. Mahdy *

Posted Date: 8 May 2026

doi: 10.20944/preprints202605.0525.v1

Keywords: quantum machine learning; hybrid quantum-classical neural networks (HQCNN); quantum convolution neural networks (QCNN); oral diseases detection; parameterized quantum circuits; principal component analysis (PCA); angle encoding

Article

Quantum-Enhanced Oral Disease Detection Using Hybrid Quantum-Classical Neural Networks

Md. Shakhawat Hossain [1,2], Md. Mehedi Hasan [2,3], Mohammad Junayed Hasan [2,4] and M. R. C. Mahdy [1,*]

[1] Department of Electrical and Computer Engineering, North South University, Dhaka, Bangladesh
[2] Mahdy Research Academy, Dhaka, Bangladesh
[3] Department of Applied Mathematics, Noakhali Science and Technology University, Noakhali, Bangladesh
[4] Department of Computer Science, Johns Hopkins University, Baltimore, United States
* Correspondence: mahdy.chowdhury@northsouth.edu

Abstract

Around the half of the population in the world are affected by oral diseases, making it one of the most common health conditions. Quantum implementation in medical domain has revealed its potential and versatile applicability especially in medical imaging. This paper explores oral disease identification using hybrid quantum-classical neural networks (HQCNN) and quantum convolution neural networks (QCNN). Our work investigates the possibilities of quantum machine learning in processing complicated dental image data and the contributions it can make in oral healthcare. We implemented a hybrid and a pure QNN leveraging Qiskit framework and a whole dataset of annotated oral disease dataset. Our 8 qubit structured QCNN model and 2 qubit architecture of HQCNN model extract the image features by encoding the features into quantum circuits enabling more expressive demonstration employing fewer parameters. The final result showcases that QCNN and HQCNN perform better than CNNs in disease classification and promise better accuracy, generalization and computational efficiency. This experiment highlights a pioneering step in applying quantum inspired models for oral diagnostics, identifying promising avenues for improving oral healthcare worldwide.

Keywords: quantum machine learning; hybrid quantum-classical neural networks (HQCNN); quantum convolution neural networks (QCNN); oral diseases detection; parameterized quantum circuits; principal component analysis (PCA); angle encoding

---

1. Introduction

Oral diseases have emerged as a significant public health concern worldwide. According to the WHO Global Health Status Report, 3.7 billion people are in some way affected by oral disease and suffer from different oral conditions, the most common ailment includes dental caries (tooth decay) [1]. Oral health conditions such as periodontal diseases, oral cancers, tooth loss, and caries can be prevented by taking appropriate treatment in their early stages [1]. In the last 30 years, the cases of oral disease have increased by 50% [2]. This concerning prevalence of the main oral diseases continues to increase globally due to growing urbanization, overexposure to fluoride, limited access to medical care, a high sugar diet, and the consumption of tobacco and alcohol [1]. This increasing burden leads to disfigurement, causes pain and increased mortality. Immediate and proper diagnosis is crucial to get improved outcomes. Recent development of artificial intelligence and different AI methods especially deep learning have been proved to be a powerful and useful tools in medical imaging and diagnostics. These methods can analyze complex images rapidly and more accurately than traditional methods [3]. For example, AI-based medical image analysis has shown better results in enhanced interpretation and improved abnormalities detection leading to overall better outcomes in patient treatment [3]. These results suggest great potential of AI application in oral healthcare, where automated diagnosis and

 doi:10.20944/preprints202605.0525.v1

detection of cavities and legions could improve the augmentation of clinical screening and help to accelerate earlier intervention.

Modern methods in oral healthcare extensively use machine learning and computer vision. Deep convolutional neural networks (CNNs) have been successfully applied to classification of periodontal conditions, tooth numbering and caries detection [3]. For instance, CNNs trained on large labeled datasets of intraoral photos or panoramic X-rays were able to identify teeth and diagnose oral diseases with increased accuracy [3]. Park et al. (2023) illustrated a CNN using color intraoral images to classify periodontal disease and achieved a higher accuracy than a standard ResNet [3]. However, these classical models come with notable drawbacks. They require millions of parameters and their performance and generalization significantly depend on the provision of large annotated datasets. With limited data deep models tend to show issues such as overfitting and poor performance [3]. Moreover, image with high dimensionality leads to huge computational overhead. Typical method of capturing dental images like radiographs often contain overlapping anatomical structures and noise which found to confound classical image analysis [3].

Hybrid quantum architectures can efficiently solve these limitations. The ability of encoding data into high-dimensional Hilbert spaces and exploitation of quantum entanglement and quantum superposition offer potential to extract more complex feature correlations using fewer model parameters [4,5]. For instance, Chen et al. (2022) demonstrated a quantum convolutional neural network (QCNN) trained for a problem related to high-energy physics classification and showed results of faster learning and higher accuracy compared to an equivalent classical CNN [4]. A recent study reported by Li et al. (2025) showed that a distributed hybrid quantum-classical CNN employing quantum circuit splitting successfully extracted complex features with fewer parameters and qubits, producing excellent performance while using reduced resources [5].

In medical imaging, hybrid quantum classical structures are already showing greater advantages. Most recent research demonstrates that quantum-augmented models are parallel or better in performance when compared with classical models. For instance, a hybrid quantum classical CNN was found to achieve an accuracy of >97% on brain MRI classification tasks and quantum transfer learning method proved to outperform classical baselines in COVID-19 X-ray classification with improved accuracy and time [6]. These studies illustrate that quantum structures can boost feature learning and can be generalized in medical contexts. To our awareness, however, no prior work has integrated quantum or hybrid quantum classical structures to detect or diagnose oral diseases. The success of quantum inspired models in other medical imaging domains encourages exploring the similar architectures in dental diagnostics with an aim to bypass the data and complexity issues of classical methods.

We present three different models in this paper to classify prevalent oral diseases: a traditional CNN, a purely quantum CNN (QCNN) and a hybrid quantum-classical convolutional neural network (HQCNN). We leverage the Qiskit framework to implement these models. We curated our dataset from a Kaggle collection of oral diseases. It features annotated images depicting dental caries, gingivitis, calculus, discolorations, ulcers among other oral conditions. We selected only the Calculus and Hypodontia as our oral disease dataset for binary classification. We prepare the images for training by performing quality control, followed by size standardization and normalization. We then broaden our training data by incorporating classical augmentations such as flips, rotations, etc. Within the HQCNN architecture, parameterized quantum circuits are combined with classical convolutional layers. The workflow involves extracting convolutional feature maps from the input image and encode them into a quantum circuit. The output measurements are then fed into subsequent classical layers. Functioning as a quantum analogue of a CNN, the QCNN transform image patches into qubit states and then processed using quantum convolution and pooling operations, as established in prior work [5]. The CNN architecture leverages a hierarchical structure of classical convolutional layers and extracts spatial features which are then directed through nonlinear activation functions and pooling layers.

 doi:10.20944/preprints202605.0525.v1

The models are trained via hybrid optimization routines: we combine the use of Qiskit's quantum circuit primitives with PyTorch to optimize parameters using gradient descent. By including quantum layers, the networks can take advantage of high-dimensional Hilbert space representations of image data and learn more expressive features using fewer parameters [5,7]. This strategy directly overcomes the drawbacks of the classical approach: the feature extractor in the quantum framework can enhance the generalizability on small medical images, and the mixed design decreases the overall computational load. We compare HQCNN and QCNN to a classical CNN baseline using our dataset of oral images based on their classification accuracy, efficiency in training, and resistance to overfitting.

This paper explores the successful application of Hybrid Quantum-Classical Neural Networks (HQCNN) and Quantum Convolution Neural Networks (QCNN) in oral disease detection by creating the structure of an efficient HQCNN and QCNN model and comparing with typical CNN to transform the diagnostic landscape. This step in the quantum realm can be a pivotal step toward a paradigm shift for global oral healthcare.

We can distill our principal contributions into the following points:

1. **Novel HQCNN application to oral diagnosis:** We developed a hybrid quantum-classical CNN for oral disease classification, thereby establishing one of the nascent applications of quantum-inspired models within the field of Oral diagnostics.
2. **Comparative Analysis of CNN, QCNN, and HQCNN:** Our research utilize real-world clinical image data and systematically evaluate classical, fully quantum (QCNN), and hybrid quantum-classical (HQCNN) CNN models. This comparative analysis elucidates the performance and benefits of quantum layers, illustrating how quantum integration can improve accuracy and generalization over purely classical models.
3. **Feasibility of quantum-enhanced dental AI:** Our Qiskit implementation of HQCNN and QCNN models strongly demonstrate the viability of these models and provide empirical proof of their capacity to advance diagnostic accuracy and efficiency in dental diagnosis.

The paper is organized as follows: Section 2 provides a examination of existing literature in quantum and hybrid neural networks relevant to medical imaging. Section 3 delineates the dataset, model architectures, and our implementation methodology. Our experimental results and their analysis can be found in section 4. Section 5 offers a discussion of the significance of the findings and their broader implications. Finally, Section 6 concludes the paper by summarizing the paper and outlines future research directions.

2. Related Work

2.1. Challenges in Hybrid Quantum-Classical CNNs for Medical Imaging

Quantum machine learning (QML) is proved to be able to enhance pattern detection by utilizing high-dimensional quantum feature spaces, yet practical implementation in medical imaging confronts substantial barriers. HQCNNs are met with limited qubit numbers, noisy intermediate-scale quantum (NISQ) hardware and circuit design limitations [8]. Some optimization problems and training difficulties including sensitivity to noise and barren plateaus make it more challenging for variational quantum circuits to handle complex clinical data [8]. Moreover, passing large medical image through quantum states can be computationally intensive and deploying quantum layers with classical CNNs adds complexity [8]. Recent systematic studies observe that quantum models on medical datasets rarely exhibit consistent advantage relative to classical approaches under real life circumstances [9,10]. One comprehensive analysis, for instance, reported no consistent trend or quantum advantage in health data tasks, indicating additional overhead by quantum models without clear gains in performance [9]. These findings highlight fundamental challenges in data encoding, scalability and hybrid parameter optimization that currently restraining the empirical benefit from QML in healthcare AI.

 doi:10.20944/preprints202605.0525.v1

2.2. Hybrid Quantum-Classical CNNs in Medical Diagnostics

In spite of these limitations, several recent works have started to apply hybrid QML architectures to image classification in medical domain. Xiang et al. (2024) evaluated a classical CNN integrated with quantum convolutional layer for breast cancer diagnosis [11]. Their proposed quantum-classical CNN utilizes parameterized quantum circuits to learn and extract features from dataset. On various breast cancer datasets, their developed QCCNN has achieved higher accuracy than a classical CNN, indicating improved generalization through quantum enhanced network. Ajlouni et al. (2023) adopted a similar approach and introduced a HQCNN for brain tumor prediction [12]. They tested the HQCNN model on a large dataset of brain tumor from Kaggle using classical convolutional layer combined with a quantum circuit as well as an adaptive optimizer. The CNN baseline was outperformed by the hybrid model with the help of quantum processing [12]. Senokosov et al. utilize two QML models for imaging. They merge hybrid quantum neural network into parallel quantum circuits and achieved a high accuracy of 99.21% and 99% on the MNIST dataset and Medical MNIST dataset respectively. This experiment showcases the potential of quantum layers in differentiating common characteristics of input data [13]. QCCNNs are further explored by other researchers for different imaging modalities. Matic et al. (2022) propose various hybrid QCCNN models for radiological image classification [14]. They applied their networks to 2D and 3D CT scan data (including images of liver and pulmonary lesions) to find that the performance of the QCCNN is similar to the classical models on these tasks [14]. This parity of performance implies even small quantum circuits can match the classical models in feature extraction of medical imaging, spurring further exploration. Building on the constrains of available hardware, Li et al. (2025) introduce a distributed HQCNN employing circuit cutting. Their method splits a convolutional circuit of 8-qubit into smaller parts implementable over just 5 qubits, there by extending QML to larger networks [5]. The distributed HQCNN model demonstrated strong binary and multi-class classification performance across three imaging datasets, exhibiting greater accuracy and reduced parameter compared to existing methods [5]. Tantawi et al. evaluated QCNN on datasets such as PneumoniaMNIST and VesselMNIST3D and were able to achieve promising results, demonstrating the potential of quantum processing in automatic medical image classification [15]. Several experiments focus on domain specific applications. Khan et al. (2024) deploy a QCNN on brain tumor detection, achieving an incredibly high test-set accuracy of 99.67% [16]. This specialized pure quantum CNN architecture proves the potential of quantum network structure in oncological imaging. On the domain of neurological diagnostic, Islam et al. (2025) introduce a hybrid architecture of classical-quantum CNN to detect Alzheimer's disease from 3D MRI scans [17]. After implementing a small 3-qubit quantum circuit along with diffusion-aided data augmentation as a method to tackle class imbalance their experiment resulted in a remarkably high accuracy of 97.50%. This result suggests quantum structure can beat the traditional classical baselines leveraging substantially reduced parameter counts [17]. Experiments on infectious diseases diagnostics have also been reported. Rao et al. (2024) design a hybrid framework merging a specialized classical convolutional neural network (CCNN) and quantum classifiers to detect respiratory lung diseases from chest X-ray images [18]. Their developed model was tested on thousands of X-ray images and the model attains testing accuracy of 98.1%, exceeding the performance of multiple state-of-the-art models such as InceptionNet, MobileNet, ResNet50 and VGG16 [18]. More recently, Asadoorian et al. (2024) propose a QCNN model incorporating transfer learning for COVID-19 CT classification [19]. Combining classical and quantum capabilities, a new hybrid model — which integrates a pre-trained VGG16 feature extractor merged with a quantum convolutional layer — demonstrated strong performance on CT scan data, reaching around 96.8% accuracy [19]. The finding indicates the complementary role of quantum layers and deep-transfer features in rapid COVID screening.

2.3. Gaps and Motivation for Oral Health Applications

The above studies collectively highlight hybrid QML's versatile applicability in medical imaging across various domains, including cancer detection, infectious disease screening and neurological disorder analysis. Interestingly, while applications of convolutional neural networks are becoming more common in dental diagnostics, none of the prior works we have discussed so far specifically focus on dental or oral conditions [20], no existing study has attempted to extend existing quantum inspired architectures to oral pathology or imaging. This represents an obvious gap: oral disease imaging (such as intraoral radiographs and dental images) involves its own set of challenges (for instance multi-label pathology and complex high-dimensional texture patterns) that have been left to explore using QML. Through our work, we are bridging this crucial gap by introducing the first hybrid quantum-classical CNN and pure quantum CNN specifically developed for oral health diagnosis. By integrating quantum layers and combining with an oral image classifier, we address the specific requirements of effectively capturing the distinct characteristics of oral disease data. This opens a new avenue for quantum-powered AI in medical field, showcasing the unprecedented ability of QCNNs and HQCNNs to enhance accuracy and efficiency in oral condition diagnosis.

3. Materials and Methods

3.1. Materials

3.1.1. Dataset

The oral disease dataset is sourced from Kaggle. We only included calculus and hypodontia as dental disorder for our dataset for binary classification. Photos taken from several hospitals and reliable dental websites make up the dataset. The dataset perfectly shows the authenticity and variety of oral diseases. Generalization was enhanced by rotation, flipping, scaling, and noise addition among data augmentation techniques. The dataset consisted of training sets of 70%, validation sets of 15%, and test sets of 15%. [21]

3.1.2. AI Models

In this paper, we have developed two AI models, HQCNN (Hybrid Quantum-Classical Neural Networks) and QCNN (Quantum Convolution Neural Networks), to classify the images from the dataset and identify the specific oral disease. Then we evaluated the two models as well as an existed CNN to determine which one performs better for the given oral disease diagnosis tasks.

Hybrid Quantum-Classical Neural Networks (HQCNN): Combines a parameterized quantum circuit with classical convolution layers to improve feature extraction [22].

Quantum Convolution Neural Networks (QCNN): Using Qiskit, encodes images into quantum states and implements quantum convolution and pooling layers [23].

Convolutional Neural Networks (CNN): An existed regularized CNN model specially developed for binary classification. The model is composed of a fully connected head following by 2 convolutional blocks.

3.1.3. Implementation Details

All the computational tasks in the paper were performed using cloud technology. We conducted the computational experiment of our model on Google Colab utilizing one single Tesla K80 GPU with 12GB GDDR5 VRAM and 2496 CUDA cores [24]. This resource is sufficient for the training and evaluation of our models. The quantum components were carried out on Qiskit's AerSimulator [25]. The dataset was downloaded from Kaggle and hosted in Google Drive. The models were built using Python with libraries including TensorFlow, PyTorch, matplotlib, Qiskit, pylatexenc, and scikit-learn.

3.2. Methods

3.2.1. Design of Hybrid Quantum-Classical Neural Networks (HQCNN)

Our proposed HQCNN model is designed to enhance binary classification on oral dataset by harnessing the power of quantum processing. It integrates a classical deep learning (particularly a fine-tuned ResNet18 backbone) merging with a well parameterized quantum layer, utilizing both classical and quantum methods to process information.

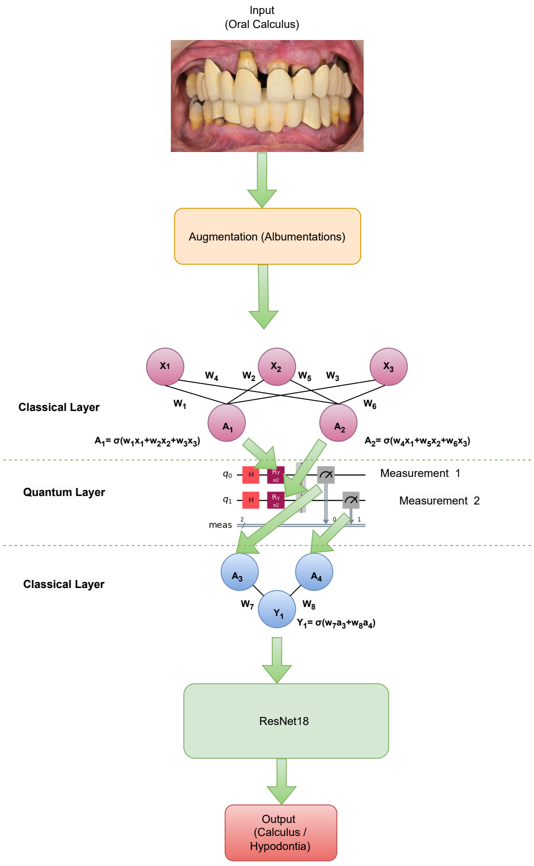

Figure 1. Overview of the structure of Hybrid Quantum-Classical Convolutional Neural Network (HQCNN). The architecture employs a classical ResNet18 working as a backbone combining with a 2-qubit parameterized quantum circuit. A trainable bias added by the quantum layer at the final logits, increasing the ability of the hybrid model to capture complex features and patterns.

Data Preprocessing and Augmentation

After resizing the images to 150×150 the augmentation was completed through horizontal flipping, random rotation, brightness/contrast adjustment and Gaussian blurring. Then channel wise normalization was performed employing a mean and standard deviation of 0.5. For loading the data into the model and transforming them systematically, a custom PyTorch Dataset class was used.

Classical Backbone

Using ImageNet we train a ResNet18 and employ it as the classical backbone of our HQCNN model. Except for layer4 and the final connected layer, all the other convolutional layers are remain fixed. The final fully connected ResNet layer is substituted with a linear projection that map to a 2D latent space, matching the number of qubits in our setup.

Quantum Layer

A quantum circuit of 2-qubit is constructed utilizing Ry rotations followed by a Controlled-NOT (CNOT) gate and full measurement. The selection of these elements for constructing quantum circuit is not arbitrary rather they work as foundational building blocks for quantum computation. Ry gates are important single qubit gates which allow rotation qubit around Y-axis by predefined angle θ [26]. This is crucial for fixing the qubits in specific superposition states, also enabling implementation of different unitary transformations required for algorithm execution[26]. The rapid parameterization of this Ry gate gives us control over the vector state of the qubit enabling the transpilation of classical

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 May 2026 doi:10.20944/preprints202605.0525.v1

data into quantum states [26]. The Ry gate is demonstrated by a 2×2 unitary matrix that works on the state vector of a single qubit.

$$RY(\theta) = \exp\left(-i\frac{\theta}{2}Y\right) = \begin{pmatrix} \cos\left(\frac{\theta}{2}\right) & -\sin\left(\frac{\theta}{2}\right) \\ \sin\left(\frac{\theta}{2}\right) & \cos\left(\frac{\theta}{2}\right) \end{pmatrix}$$

CNOT gate is a quantum logic gate that works on two qubit, one is designated control qubit and the other one is called target qubit. This gate actually operate as Pauli-X (NOT) on target qubit which means it performs action if and only if the state of the control qubit is $|1\rangle$[27]. CNOT gate is important because this gate is able to do primary entangling when applied to quantum qubits[27]. This enable many advanced quantum computational benefits over classical protocols. We can demonstrate the CNOT gate by a 4×4 unitary matrix that works on 2-qubit's joint state vector.

$$CNOT = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

The output of the CNOT gate is then calculated by measurement operation which perform its action on both qubits, serving as an important interface between the quantum and classical states. The circuit setup is illustrated below:

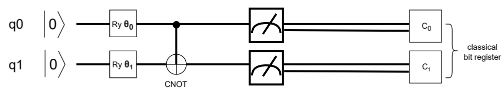

Figure 2. 2-Qubit Quantum Circuit with Ry gate, CNOT gate and measurement; calculating quantum computation and passing to classical state.

The architecture helps to learn the circuit parameters durin training. Gradients are measured through parameter shift rule enabling backpropagation in the quantum layer. The quantum layer is integrated through a custom autograd function built in a PyTorch module.

Classification Head

Once produced, the output of the quantum layer undergoes significant reshaping before being integrated with the classical features, establishing a richer quantum-augmented representation. This enhanced representation is then fed into a feedforward network comprise of two layers, which ncorporate ReLU activation functions. The function ultimately computes the final logits for the classification task.

Training Strategy

Supported by the technique of class-balanced sampling, our training regimen characterizes 5-fold stratified cross-validation. Class-balanced sampling method addresses possible disparities and ensures fair representation. A function called weighted cross-entropy loss is incorporated to handle the class imbalance effect. Optimization process is executed by Adam algorithm, coupled with a ReduceLROnPlateau (a learning rate scheduler). The training of the model proceeds by spanning 10 epochs per fold (total 5 folds). Furthermore, we applied the method of gradient clipping to effectively counter the exploding gradients.

Evaluation and Interpretability

The performance of the model is evaluated by measuring standard performance metrics including accuracy, precision, recall, F1 score, and AUROC. Model interpretability is prioritized and significantly enhanced by applying Grad-CAM and LIME. Grad-CAM reveals spatial regions relevant for class discrimination, while LIME offers insights into important local feature attributions. Additionally, the model displays dedicated modules, illustrating how both components i.e., classical and quantum influence the overall decision-making process.

3.2.2. Design of Quantum Convolution Neural Networks (QCNN)

We developed our QCNN binary classification model by modifying the existing IBM Qiskit machine learning project. Quantum Convolution Neural Networks (QCNN) work in a mechanism that we can relate to the classical neural network. First, a pixelated image is encoded into a quantum circuit using Qiskit's ZFeatureMap; after that, it passes a sequence of alternating quantum convolution and pooling layers. Until only one qubit remains, this procedure progressively lowers the dimensionality of the circuit, which is then measured to categorize the input image. Every layer in a QCNN has parameterized circuits, which means the output of one layer depends on the changeable parameters in another. These parameters are tuned throughout training to reduce the loss function of the network [28].

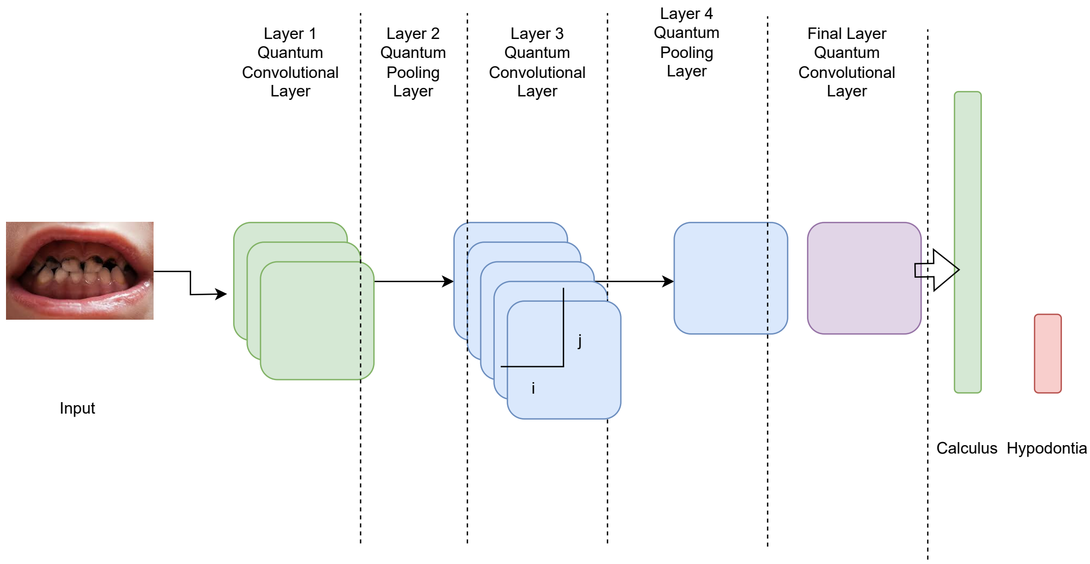

Figure 3. A QCNN model used for classifying dental images consists of several convolution and pooling layers that reduce the dimensionality of data to classify images either as calculus or hypodontia.

Building the QCNN Components

For the QCNN, convolution and pooling layers are implemented in terms of gates applied to a Quantum Circuit. Each one of these layers will have parameters; we will modify these parameters over the training process to raise the model efficiency, reduce the loss function and train the QCNN to distinguish horizontal from vertical lines. For every unitary gate operating on a pair of qubits, generators are the Gellmann Matrices. Our parameterized circuit was created on the two-qubit unitary operation. Here, Every unitary matrix in U(4) may be broken down such that [29]:

$$U = (A1 \otimes A2) \cdot N(\alpha, \beta, \gamma) \cdot (A3 \otimes A4)$$

Here, $A_j \in SU(2)$, Tensor product is denoted by $\otimes$, and

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 May 2026 doi:10.20944/preprints202605.0525.v1

$$N(\alpha, \beta, \gamma) = \exp(i[\alpha\sigma_x\sigma_x + \beta\sigma_y\sigma_y + \gamma\sigma_z\sigma_z]),$$

Here, α, β, γ are the parameters [29].

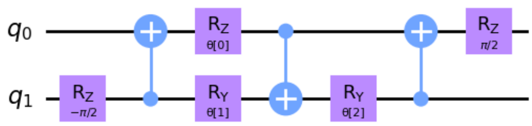

Figure 4. Parametrized two-qubit unitary circuit for $N(\alpha, \beta, \gamma)$.

Quantum Convolution Layer

We have defined the earlier established unitaries. First, we apply the two-qubit unitary to all even couples of qubits to develop a function of the convolution layer in our Quantum Convolution Neural Network (QCNN). Next, we apply it in a circular coupling manner to the odd pairs of qubits, therefore coupling neighboring qubits as well. Additionally, the first and last qubits are connected through a unitary gate [28].

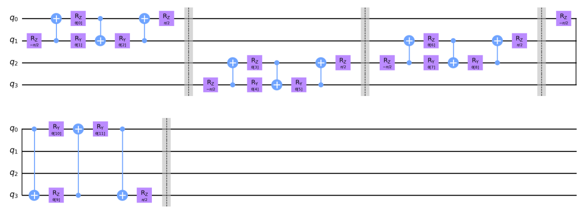

Figure 5. Gate organization in the QCNN's Convolution Layer.

Quantum Pooling Layer

The pooling layer has to be applied differently than in conventional neural networks. We first group N qubits into pairs to attain an artificial decrease. We use a generalized two-qubit unitary operation for every pair. We discard one qubit from each pair after this transition, lowering the total qubit count carried out in the network. This approach essentially combines the data from both qubits into a single qubit. The obsolete qubit is set aside while the unitary operation codes the data from one qubit into the other. We build a Quantum Convolution Neural Network (QCNN) pooling layer by applying this method that efficiently alters the dimensions of our N quantum circuit to N/2 [28].

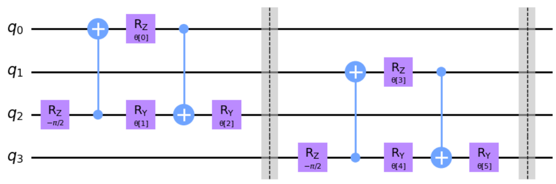

Figure 6. Gate organization in the QCNN's Pooling Layer.

QCNN Modeling

Alternating convolution and pooling layers make up our QCNN. Our QCNN will leverage eight qubits. We apply a feature map to encode our dataset into our QCNN. We have tried different feature maps on our dataset and found that the most suitable one is ZFeatureMap with highest accuracy, outperforming ZZFeatureMap, amplitude encoding and Qubit encoding. The benchmark of these methods are illustrated in table One of Qiskit's built-in ZFeatureMap allows one to generate a feature map. QCNN performs best when the Zfeature map is applied. We estimate the expected value of the Pauli Z operator of the last qubit to categorize our oral disease image dataset into vertical and horizontal lines. We can deduce from the obtained value of either +1 or -1 that the input image had a horizontal or vertical line [28].

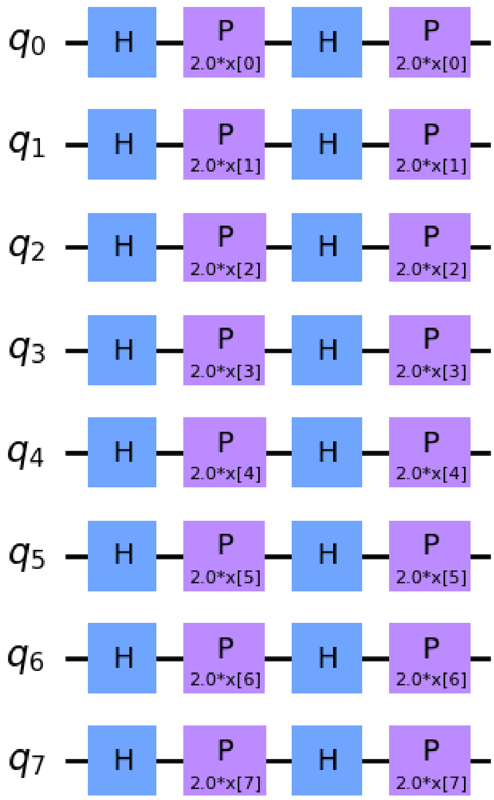

Figure 7. ZFeatureMap to Encode Oral Disease Dataset into QCNN Model.

QCNN Training

We measure from the output circuit in order to categorize our system. The value we get will thus define whether our input data comprises a vertical or a horizontal line. For the last qubit, our selected measurement is $< Z >$ (Pauli Z qubit's expectation value). We derive +1 or -1 by measuring this expectation value, corresponding to a vertical or horizontal line. A commonly utilized numerical optimization method for classification machine learning algorithms, the COBYLA optimizer, will be utilized to train our classifier. After that, we used the built-in Neural Network Classifier of Qiskit Machine Learning to train our model; we also included the callback function, optimizer, and operator for our QCNN [28].

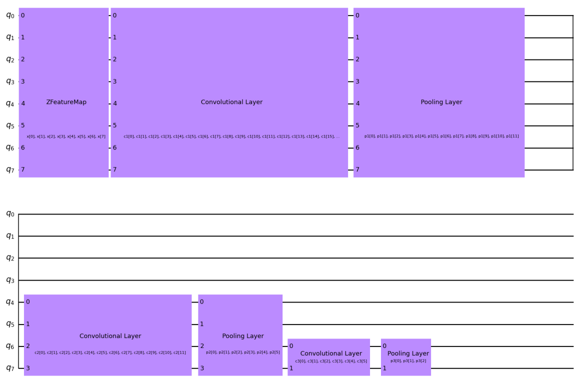

Figure 8. Reducing the dimensionality of QCNN from eight qubits to one.

3.2.3. Design of Convolutional Neural Network

We utilize an existed regularized CNN model developed for binary classification as a robust classical benchmark. We trained the model on a labeled dataset of oral diseases (calculus and hypodontia). We applied suitable preprocessing methods and augmentation strategies to improve generalization.

Data Handling and Preprocessing

We load the dataset using an API (image_dataset_from_directory) from TensorFlow. We then split the dataset in 80:20 ratio for training and validation. As a preprocessing technique we resized all the images to 32×32 pixels and further normalizing them to the [0,1] range.

Model Architecture

The architecture of the CNN model is developed with a fully connected head following by 2 convolutional blocks. The blocks contain standard regularization elements to prevent overfitting, which is very common in datasets consist of small-scale medical image. The model process the output logits employing SparseCategoricalCrossentropy(from_logits=True), which confirm analytical stability in softmax computation.

Training Strategy

The model implement Adam optimizer for training with an initial specified learning rate of 0.001. The model continue to train for a maximum of ten epochs with an integration of early stopping to stop training if there is no improvement in validation performance for 5 successive epochs. A well established learning rate scheduler called ReduceLROnPlateau is also used to half the learning rate in case the model observes a plateau in the validation loss over three epochs.

Evaluation Metrics

Performance was quantified using metrics Accuracy, Macro-averaged Precision, Recall, and F1 Score, (AUROC) and Confusion Matrix. Class probabilities were derived by computing softmax probabilities from the output logits. A visualization of missclassification of the model is depicted by a full confusion matrix.

3.2.4. Data Augmentation and Preprocessing

HQCNN binary classification model is designed to identify dental calculus and dental hypodontia images and categorize those images into two types of digits (0 or 1). Utilizing PyTorch for image ingestion with Albumentations for effective image transformation, we load the dental disease images and scale those to 150x150 pixels, then those dental disease images are normalized by :

$$x' = \frac{x}{255.0}$$

[30]

For QCNN model, the oral disease images are converted into 2 x 4 pixelated images. Every pixel is given a random value between 0 and pi/4 to produce a noisy background. Images with a vertical line of +1 and images with a horizontal line of -1 can be labeled; thus, QCNN can learn to distinguish between the two patterns. Following is the augmentation algorithm for the QCNN model [31]:

1. **Input:** Original oral disease images I
2. **Quantum Data Encoding:** Input images are encoded into quantum states $I_{\text{quantum}} = \text{ZFeatureMap}(I)$
3. **Quantum Rotations:** Applying rotations using CX, RY, and RZ gates to adjust qubit states
4. **Entanglement:** Using Entanglement between qubits to recognize the complex feature
5. **Phase Shifts:** applying phase shifts to specific qubits in order to replicate various data viewpoints
6. **Output:** Final quantum state is measured $I_{\text{final}} = \text{measure}(I_{\text{augmented}})$

3.2.5. Dimensionality Reducing and Quantum Data Encoding

In the HQCNN model we employ ResNet-18, a pretrained extractor which extract feature as a 512 dimensional feature vector after the operation of global pooling. Our HQCNN model work with 2 qubits which is much smaller than 512. The model freezes all the layer except for layer 4, so that only the linear layer can act as trainable projection from 512 $\longrightarrow$ 2 qubits. We then feed the dimensionally reduced vector into the quantum circuit through QLayer.

In the QCNN model, after transforming raw images to 16×16 pixels of grayscale images, we convert them into a 256 dimensional vector by following the process of flattening the pixel grid. We call this data matrix

$$X \in \mathbb{R}^{N\times 256},$$

N being the number of samples.

It reduces the vector to 8 Principal Components by computing the top 8 eigenvectors of the sample covariance matrix of X, establishing

$$W_8 = [w_1, w_2, \ldots, w_8] \in \mathbb{R}^{256\times 8}.$$

The projection

$$X_{\text{PCA}} \;=\; X\, W_8$$

generate an $N \times 8$ principal component score matrix for every image. Then these 8 dimensional vectors are transformed into valid quantum state for quantum encoding by normalizing each row in the matrix to unit length. Specifically, if

$$x_i \in \mathbb{R}^8$$

is the i-th reduced vector then we compute

$$x_i^{(\text{norm})} \;=\; \frac{x_i}{\|x_i\|_2}$$

So that

$$\|x_i^{(\text{norm})}\| = 1$$

These components become our required amplitude-encoded inputs for the 8-qubit feature map. The PCA helps us to reduce the dimension of our image from 256 to 8, so that it matches our 8-qubit cicuit. It also helps us to eliminate the low-variance noise present in the pixel signals while keeping only principal components.

3.2.6. Training and Hyperparameter Optimization

The HQCNN model is trained by applying the method of the multi-class cross-entropy minimization on predicted probabilities and actual labels. For a batch containing N samples and C different classes, the model calculate the cross-entropy loss function as [32]:

$$L \;=\; -\sum_{i=1}^{N}\sum_{c=1}^{C} y_{ic}\, \log\!\big(\hat{y}_{ic}\big)$$

y_{ic} is 1 for sample i belonging to class c, else 0. $\hat{y}_{ic}$ represents the softmax probability specified to class c. The Adam algorithm work as optimizer with the initial learning rate 1×10^{-3}. This helps us to Combine momentum with adaptive steps in each parameter for faster convergence in our predefined hybrid networks. We employed early stopping mechanism with 10 epochs and after completing each epoch, the validation loss is calculated. If there is no decrease in validation loss up to 10 consecutive epochs then the algorithm stop training model to avoid overfitting.

To handle the Learning-Rate Scheduling we utilize **ReduceLROnPlateau** which is triggered when there is no significant improvement for a adjustable number of epochs. It works by multiplying the current learning rate by 0.5 allowing excellent weight updates when progress is stopped. This helps the model to adapt tighter minima.

Other hyperparameters: Max epochs is set high and optimized through effective early stopping. The initial batch size is 16 which is balanced in each of the five fold through weighted sampling.

 doi:10.20944/preprints202605.0525.v1

Parameters of quantum layer contain number of shots and shift angle ($\pi/16$). 5-fold cross-validation by StratifiedKFold yields reliable performance assessments.

3.2.7. Model Evolution and Matrices

The trained QCNN and HQCNN model is assessed across several criteria:

$$\text{Accuracy} = \frac{\text{Number of Correct Predictions}}{\text{Total Samples}} \times 100\%$$

F1-score, precision, and Recall are calculated to handle multi-class categorization using macro-averaging. A one-versus-one (OVO) method computes the Area Under the Receiver Operating Characteristic curve. The confusion matrix is visualized to check misclassifications across classes [33].

4. Results and Analysis

4.1. Main Results

QCNN and HQCNN are designed especially for binary classification tasks [34]. So, We investigated QCNN and HQCNN architectures for binary oral disease identification. We conducted several experimental runs for both the HQCNN and QCNN configurations to assess our models' robustness and we also run a CNN model to compare how the quantum models perform against classical model. Early stopping was used to avoid overfitting and maximize training length, guaranteeing that every model attained its best performance without needless epochs. Table 2 presents performance metrics:

Table 1. Performance comparison (mean ± std) across 3 runs.

Model	Accuracy	Precision	Recall	F1	AUROC	Runtime (s)
HQCNN	0.997 ± 0.0008	0.997 ± 0.0008	0.997 ± 0.0008	0.997 ± 0.0008	1.000 ± 0.0000	3912.67 ± 187.62
QCNN	0.714 ± 0.0074	0.745 ± 0.0068	0.694 ± 0.0022	0.690 ± 0.0034	0.694 ± 0.0021	2160.00 ± 48.99
CNN	0.698 ± 0.0342	0.610 ± 0.2018	0.581 ± 0.0631	0.540 ± 0.1105	0.689 ± 0.1351	1883.00 ± 32.03

Table 2. Paired t-test statistics between models (p-values).

Metric	HQCNN vs QCNN	HQCNN vs CNN	QCNN vs CNN
Accuracy	2.68×10^{-4}	0.0063	0.5256
Precision	3.93×10^{-4}	0.1127	0.4528
Recall	1.13×10^{-5}	0.0111	0.1216
F1 Score	3.71×10^{-5}	0.0278	0.1875
AUROC	2.42×10^{-5}	0.0827	0.9632

Our HQCNN 2-qubit model is the best performing model achieving accuracy 99.8% in the binary classification task of oral disease detection. Alternatively, the QCNN 8 qubit model registered a lower accuracy of 72.4% compared with the HQCNN model in the binary oral disease classification task.

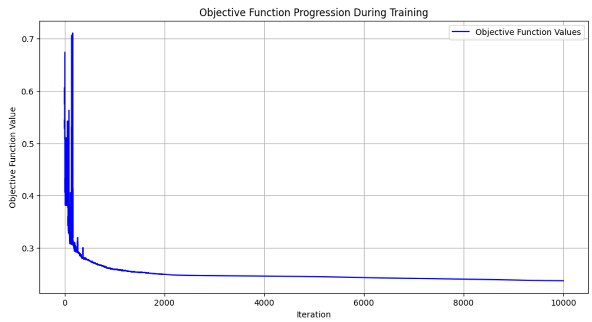

Figure 9. Objective function progression during training.

4.2. Analysis of Results

With a top accuracy of 99.8%, the HQCNN model performed noticeably better in the binary classification task than the QCNN model and the CNN model. This demonstrates the value of combining the traditional convolution layer with the quantum layer, adding a definite advantage in more straightforward tasks involving binary classification, and proves that the hybrid quantum-classical model captures essential key features more effectively than the pure quantum convolution model and typical classical model.

In Figure 10, the boxes delineate the interquartile range (IQR), with the embedded horizontal line denoting the median. The whiskers extend to show the full data span and any visible outliers revealing fluctuations in performance. The Figure also shows that HQCNN leads in accuracy across all test runs indicating minimal variance and high consistency. QCNN maintains a moderate and consistent accuracy while CNN demonstrate volatile performance. The HQCNN model has proved its reliability by maintaining a consistency in precision around 1.0. QCNN also shows slightly lower but consistent result in precision while CNN points to instability with high variance.

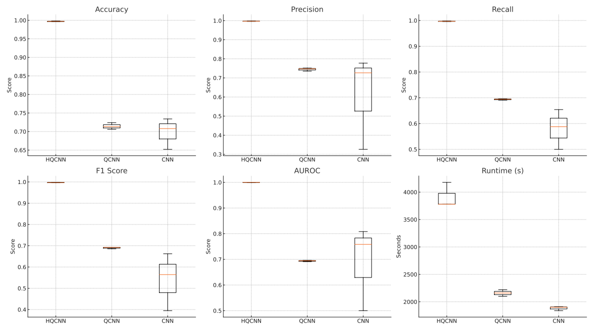

Figure 10. Performance comparison of HQCNN, QCNN and CNN with box and whisker plot across three experimental runs for each model.

4.3. Comparison with Existing Methods

In addition to traditional deep learning, salivary diagnostics, biosensing technologies, and AI-based software that uses clinical and histopathological data are other approaches currently used to detect oral diseases [35]. For example, one study on salivary tests for the detection of periodontal disease assessed markers such as lactoferrin and alkaline phosphatase Oral Diagnostic Methods, and another created a mobile application [36] to detect oral disease that preserves privacy by utilizing federated learning. Although these techniques are novel, they do not directly match HQCNN regarding image-based classification performance; however, they demonstrate the variety of methods available in the field.

Regarding AI methods, hybrid strategies that combine CNNs with SVM or Decision Trees have demonstrated promise for the high-precision detection of periodontal disease Frontiers in Medical Technology [37]. However, because tasks and datasets differ, it is not easy to directly compare with HQCNN, and specific accuracy metrics for these methods vary.

4.4. Ablation Studies

4.4.1. Ablation on HQCNN

The impact of significant architectural and hyper-parameter choices on the performance of the hybrid quantum-classical model was examined. The studies mainly focused on quantum circuit depth, encoding technique, and PCA-based dimensionality reduction.

Using an entirely classical deep learning baseline, the hybrid quantum-classical model was compared to show that, with fewer parameters, the quantum-enhanced technique obtained equivalent accuracy, implying possible advantages in efficient learning. Still, completely classical models stayed stronger against noise and required less specialized gear.

Our HQCNN model contain a single quantum layer which is implemented by 2-qubit circuit. It is observed that using one single quantum layer reduce circuit complexity, increase efficiency and improve the feasibility of the model for transpilation on actual quantum hardware [38].

4.4.2. Ablation on QCNN

Various quantum encoding techniques were investigated, including PCA-based angle encoding, ZFeatureMap, ZZFeatureMap, amplitude encoding, Qubit encoding, pooling and cropping (see the benchmark results in Table 3). We found that PCA-based angle encoding and ZFeatureMap demonstrated the greatest steady and accurate classification performance for the QCNN.

Table 3. Test Accuracy of QCNN Model with Various Feature Representations (Mean ± Std. Dev.).

Configuration	Accuracy (%)
PCA	71.45 ± 0.74
ZFeatureMap	71.45 ± 0.74
Amplitude Encoding	51.57 ± 2.31
Qubit Encoding	59.93 ± 0.35
ZZFeatureMap	58.83 ± 3.06
Quantum Pooling	64.36 ± 0.63
Cropping	66.07 ± 1.32

These results underline the need to carefully select quantum circuit depth, feature encoding techniques, and data pretreatment methods to obtain an optimal balance between performance and computing practicality.

4.5. Interpretability

The GradCAM heat map overlay provides data on which parts of the image most influence the forecast of the HQCNN. Area in red (visible in Figure 11) indicate powerful positive effect on the prediction. In our oral imagery, Grad-CAM's interpretability allows it to prominently pinpoint calculus or the locations where teeth are missing.

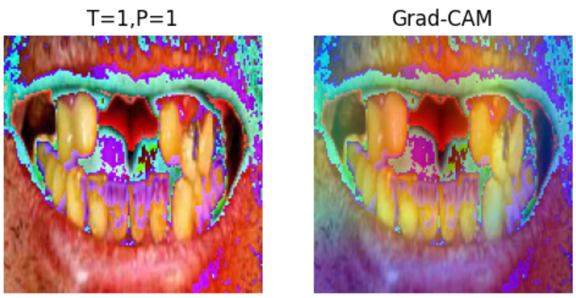

Figure 11. GradCAM heat map Overlay.

The act of visualizing the network's attention, provides a clear indication when it gravitates towards irrelevant features (e.g., background), guiding adjustments to data augmentation, training procedures and architectural design. This increases the reliability of the model. Showing heatmaps In healthcare applications serve as a crucial bridge over the "black-box" gap, enabling dentists to visually verify if the model's predictions are consistent with actual disease signs. Leveraging Grad-CAM within our hybrid model, we ensure that, even with the addition of a quantum layer, the model provides a clear insight into the decision making process of classical feature extractor which increases clarity for research and real world scenarios.

4.6. Error Analysis

Figure 12 illustrates the confusion matrix for HQCNN model.

Code available at: https://github.com/shakhawathossain07/QUANTUM-ENHANCED-ORAL-DISEASE-DETECTION-USING-HYBRID-QUANTUM-CLASSICAL-NEURAL-NETWORKS

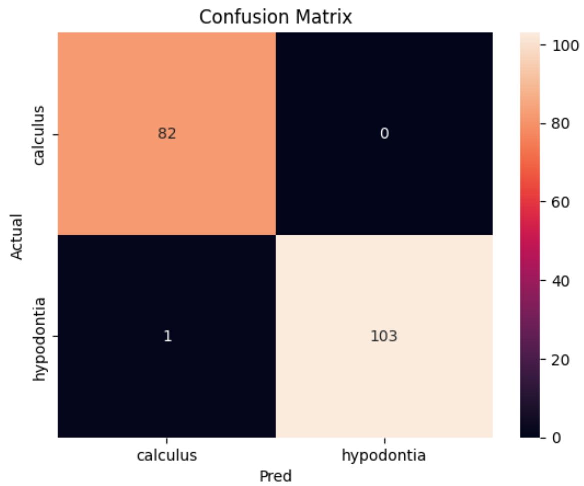

Figure 12. Confusion Matrix for HQCNN.

For both classes, the HQCNN model exhibits impressive precision and recall. Low misclassification (only one) highlights excellent generalization of the model. This indicates the model is extremely effective differentiating between the two oral conditions, delivering almost flawless binary classification.

Figure 13 visualizes the confusion matrix for QCNN which reveals several unambiguous misclassification trends for the QCNN model. Although the performance of the model in identifying the class 'hypodontia' is high, the model struggles detecting the class 'calculus' (with 72 misclassifications). The high recall for hypodontia suggests that the model can precisely identify most of the hypodontia cases. However the precision for hypodontia is comparatively lower because of the too many wrong detection of the calculus samples.

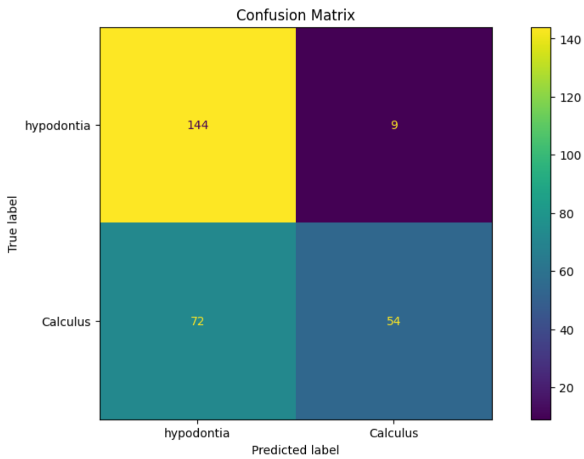

Figure 13. Confusion Matrix for QCNN.

5. Discussion

Our results expand upon past studies on hybrid quantum-classical neural networks, quantum Convolution neural networks, and quantum-enhanced diagnostic systems. Similar research has indicated that integrating quantum layers into Convolution designs can enhance feature extraction and classification accuracy in medical imaging applications [39]. The better performance of the hybrid model over the solely quantum approach in our work is compatible with earlier studies highlighting the advantages of a hybrid design in reducing quantum noise and improving learning capacity [40].

The enhanced diagnostic accuracy—99.8% in binary classification of oral disease showcases the possibility of our method for practical use in oral disease detection. Such quantum-enhanced models can provide fast and accurate diagnostic tools for early oral disease identification, assisting in minimizing diagnosis mistakes and enabling better treatment planning. Moreover, integrating hybrid quantum-classical neural networks should open the path for more resource-efficient diagnostic systems, addressing the current boundaries associated with traditional computational power in detecting medical imaging [41].

Our work is limited by present quantum hardware constraints and simulation-based evaluations, regardless of positive findings. Particularly quantum noise, limited qubit counts, and limited circuit depths provide difficulties for scaling the models. Future studies ought to concentrate on:

1. The development of more noise-resilient quantum circuits and error-reducing techniques.
2. Investigating deeper hybrid designs and better quantum encoding methods.
3. Multi-class oral disease identification.
4. Moving from qiskit Aer simulations to tests on actual quantum hardware or using quantum cloud systems.

Fully utilizing the advantages of quantum machine learning for oral disease diagnostics depends on addressing these problems.

6. Conclusions

In the last few years, machine learning and deep learning have enriched the medical industry. Different research has been conducted on disease diagnosis using machine learning and deep learning techniques. On the other hand, quantum machine learning added a new domain in the medical image classification field, which is of increasing interest among researchers. Owing to this field of quantum revolution, we developed two quantum machine learning models that have generalizability and can be utilized as a convenient, quicker, and non-invasive technique for medical staff and healthcare systems to identify oral diseases. Parameterized quantum circuits allow the recognition of improved patterns in complex datasets. The study results demonstrate that the hybrid quantum-classical neural network models can learn from augmented data effectively and achieve competitive performance in identifying different states of dental disease. This work shows the possibilities of hybrid quantum-classical neural networks for oral condition diagnosis, surpassing pure quantum convolution neural network models in the binary classification of oral disease identification problems. The suggested HQCNN model shows a viable avenue for increasing diagnostic accuracy and efficiency by using quantum feature extraction and merging traditional deep learning methods. This experimental study helps us understand the true potential of quantum-enhanced machine learning architectures. It establishes the groundwork for future experimental research integrating quantum machine learning methods with classical deep learning frameworks. Further experimental investigations are required to make our models more efficient and applicable in real-life scenarios. Future research on quantum machine learning will focus on optimizing and improving quantum circuit architectures, investigating more error reduction techniques, and exploring more quantum feature maps. Finally, it is noticeable that quantum-enhanced model architectures have the potential to improve the traditional oral disease diagnostics and provide a more secure, precise and efficient oral condition detection.

References

1. World Health Organization. (2025, Jul.) Oral health. World Health Organization. Accessed: 6 Jul 2025. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/oral-health
2. World Health Organization. Global Oral Health Status Report: Towards Universal Health Coverage for Oral Health by 2030. World Health Organization Report, 2022. Accessed: 6 Jul 2025.
3. Pinto-Coelho, L. How Artificial Intelligence Is Shaping Medical Imaging Technology: A Survey of Innovations and Applications. *Bioengineering* **2023**, *10*, 1435. Accessed: 6 Jul 2025, https://doi.org/10.3390/bioengineering10121435.
4. names not available, A. Quantum Convolutional Neural Networks for High Energy Physics Data. *Physical Review Research* **2022**, *4*, –. https://doi.org/10.1103/PhysRevResearch.4.013231.
5. Li, Y.; Qi, Z.; Li, Y.; Yang, H.; Shang, R.; Jiao, L. A Distributed Hybrid Quantum Convolutional Neural Network for Medical Image Classification. arXiv preprint arXiv:2501.06225v1, 2025. Submitted 7 Jan 2025.

6. Yan, F.; Huang, H.; Pedrycz, W.; Hirota, K. Review of Medical Image Processing Using Quantum-Enabled Algorithms. *Artificial Intelligence Review* **2024**, *57*, 300. Open access; Accessed: 6 Jul 2025, https://doi.org/10.1007/s10462-024-10932-x.
7. Chen, S.Y.; Wei, T.; Zhang, C.; Yu, H.; Yoo, S. Quantum Convolutional Neural Networks for High Energy Physics Data Analysis. *Physical Review Research* **2022**, *4*. https://doi.org/10.1103/PhysRevResearch.4.013231.
8. Shahriyar, M.F.; Tanbhir, G. Advancements and Challenges in Quantum Machine Learning for Medical Image Classification: A Comprehensive Review. In Proceedings of the Proceedings of the 2025 3rd International Conference on Intelligent Systems, Advanced Computing and Communication (ISACC), Silchar, India, 2025. Added to IEEE Xplore on 22 Apr 2025, https://doi.org/10.1109/ISACC65211.2025.10969166.
9. Gupta, R.S.; Wood, C.E.; Engstrom, T.; Pole, J.D.; Shrapnel, S. A systematic review of quantum machine learning for digital health. *npj Digital Medicine* **2025**, *8*. Published: 2 May 2025, https://doi.org/10.1038/s41746-025-01597-z.
10. Mathur, N.; Landman, J.; Li, Y.Y.; Strahm, M.; Kazdaghli, S.; Prakash, A.; Kerenidis, I. Medical image classification via quantum neural networks. arXiv preprint arXiv:2109.01831, 2021. Version 2 revised on 23 Dec 2022.
11. Xiang, Q.; Li, D.; Hu, Z.; Yuan, Y.; Sun, Y.; Zhu, Y.; Fu, Y.; Jiang, Y.; Hua, X. Quantum-classical hybrid convolutional neural networks for breast cancer diagnosis. *Scientific Reports* **2024**, *14*. Published: 21 October 2024, https://doi.org/10.1038/s41598-024-74778-7.
12. Ajlouni, N.; Özyavaş, A.; Takaoğlu, M.; Takaoğlu, F.; Ajlouni, F. Medical image diagnosis based on adaptive Hybrid Quantum CNN. *BMC Medical Imaging* **2023**, *23*. Published: 14 September 2023, https://doi.org/10.1186/s12880-023-01084-5.
13. Senokosov, A.; Sedykh, A.; Sagingalieva, A.; Kyriacou, B.; Melnikov, A. Quantum machine learning for image classification. *Machine Learning: Science and Technology* **2024**, *5*, 015040. Published: 8 March 2024, https://doi.org/10.1088/2632-2153/ad2aef.
14. Matic, A.; Monnet, M.; Lorenz, J.M.; Schachtner, B.; Messerer, T. Quantum-classical convolutional neural networks in radiological image classification. *arXiv preprint arXiv:2204.12390* **2022**. Accepted by IEEE for publication (QCE22).
15. Li, Y.; Qi, Z.; Li, Y.; Yang, H.; Shang, R.; Jiao, L. Automatic Classification of Medical Image Modality Using Quantum Convolutional Neural Network. In *Deep Learning in Medical Signal and Image Processing*; Aamir, M., Ed.; IGI Global, 2023; chapter 12, pp. 1–20. https://doi.org/10.4018/978-1-7998-8160-2.ch012.
16. Khan, M.A.Z.; Galib, A.A.O.; Innan, N.; Bennai, M. Brain Tumor Diagnosis Using Quantum Convolutional Neural Networks. arXiv preprint arXiv:2401.15804v3, 2024. Revised 25 May 2025; includes hybrid QCNN with 99.16% training and 91.47% validation accuracy across varied conditions.
17. Islam, M.; Hasan, M.J.; Mahdy, M. R.C. CQ-CNN: A Hybrid Classical-Quantum Convolutional Neural Network for Alzheimer's Disease Detection Using Diffusion-Generated and U-Net Segmented 3D MRI. arXiv preprint arXiv:2503.02345, 2025. Submitted 4 Mar 2025; achieves 97.5 % accuracy in fewer epochs than classical models:contentReference[oaicite:1]index=1.
18. Ajlouni, N.; Özyavaş, A.; Takaoğlu, M.; Takaoğlu, F.; Ajlouni, F. Hybrid Framework for Respiratory Lung Diseases Detection Based on Classical CNN and Quantum Classifiers from Chest X-Rays. *Biomedical Signal Processing and Control* **2023**, *85*, 105567. Published online Nov 2023, https://doi.org/10.1016/j.bspc.2023.105567.
19. Asadoorian, N.; Yaraghi, S.; Tahmasian, A. Pre-trained Quantum Convolutional Neural Network for COVID-19 Disease Classification Using Computed Tomography Images. *PeerJ Computer Science* **2024**, *10*, e2343. Published: 18 October 2024, https://doi.org/10.7717/peerj-cs.2343.
20. Schwendicke, F.; Golla, T.; Dreher, M.; Krois, J. Convolutional Neural Networks for Dental Image Diagnostics: A Scoping Review. *Journal of Dentistry* **2019**, *91*, 103226. PMID: 31704386; "CNNs are increasingly employed for dental image diagnostics … their usefulness, safety and generalizability should be demonstrated …" :contentReference[oaicite:1]index=1, https://doi.org/10.1016/j.jdent.2019.103226.
21. Sajid, S. Oral Diseases Dataset, 2023. Accessed: Feb. 23, 2025.
22. Liu, J.; Lim, K.H.; Wood, K.L.; Huang, W.; Guo, C.; Huang, H.L. Hybrid Quantum-Classical Convolutional Neural Networks. *arXiv preprint arXiv:1911.02998* **2019**. Accessed: Feb. 25, 2025.
23. Cong, I.; Choi, S.; Lukin, M.D. Quantum Convolutional Neural Networks. *arXiv preprint arXiv:1810.03787* **2019**. Accessed: Feb. 23, 2025.
24. Google Research. Google Colaboratory Frequently Asked Questions, 2025. Accessed: Feb. 26, 2025.
25. IBM Quantum. Exact Simulation with Qiskit SDK Primitives, 2025. Accessed: Feb. 26, 2025.

 doi:10.20944/preprints202605.0525.v1

26. IBM Quantum and Community. (2025) Rygate — qiskit circuit library documentation. Accessed: 8 Jul 2025. [Online]. Available: https://quantum.cloud.ibm.com/docs/en/api/qiskit/qiskit.circuit.library.RYGate
27. M. Ivezic. (2020, Mar.) The controlled-not (cnot) gate in quantum computing. Accessed: 8 Jul 2025; "The CNOT gate is to quantum circuits what the XOR gate is to classical circuits. . . " :contentReference[oaicite:1]index=1. [Online]. Available: https://postquantum.com/quantum-computing/cnot-gate-quantum/
28. Qiskit Community. The Quantum Convolutional Neural Network, 2023. Accessed: Feb. 26, 2025.
29. Vatan, F.; Williams, C. Optimal quantum circuits for general two-qubit gates. *Phys. Rev. A* **2004**, *69*, 032315. Accessed: Mar. 7, 2025, https://doi.org/10.1103/PhysRevA.69.032315.
30. Francois Chollet, Mark Omernick. Working with Preprocessing Layers, 2023. Accessed: Feb. 26, 2025.
31. Zhouli, L.; Wang, P.; Parampalli, U. Impact of Data Augmentation on QCNNs. *arXiv preprint arXiv:2312.00358* **2023**. Accessed: Mar. 22, 2025.
32. Zhang, Z.; Sabuncu, M.R. Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels. *arXiv preprint arXiv:1805.07836* **2018**. Accessed: Feb. 26, 2025.
33. Google Developers. Classification: Accuracy, Precision, and Recall, 2023. Accessed: Feb. 26, 2025.
34. Hafeez, M.A.; Munir, A.; Ullah, H. H-QNN: A Hybrid Quantum–Classical Neural Network for Improved Binary Image Classification. *AI* **2024**, *5*, 1462–1481. Accessed: Mar. 7, 2025, https://doi.org/10.3390/ai5030070.
35. Ramenzoni, L.L.; Lehner, M.P.; Kaufmann, M.E.; Wiedemeier, D.; Attin, T.; Schmidlin, P.R. Oral Diagnostic Methods for the Detection of Periodontal Disease. *Diagnostics* **2021**, *11*, 571. https://doi.org/10.3390/diagnostics11030571.
36. Shankara Narayanan V, Sneha Varsha M, S.A.A.G.J. Mobile Application for Oral Disease Detection using Federated Learning, 2023. Accessed: 2025-03-24.
37. Jundaeng, J.; Chamchong, M.; Nithikathkul, C. Artificial intelligence-powered innovations in periodontal diagnosis: a new era in dental healthcare. *Frontiers in Medical Technology* **2025**, *6*. https://doi.org/10.3389/fmedt.2024.1469852.
38. Group, A.R. Post-variational Classical–Quantum Transfer Learning for Medical Image Classification. *Scientific Reports* **2025**. Published approx. 5 days ago as part of hybrid CQTL framework :contentReference[oaicite:1]index=1, https://doi.org/10.1038/s41598-025-08887-2.
39. Arthur, D.; Date, P. A Hybrid Quantum-Classical Neural Network Architecture for Binary Classification. *arXiv preprint arXiv:2201.01820* **2022**. Accessed: Mar. 7, 2025.
40. Laboratory, O.R.N. Hybrid Quantum-Classical Neural Networks, 2023. Accessed: Mar. 7, 2025.
41. Zayed, S.O.; Abd-Rabou, R.Y.M.; Abdelhameed, G.M.; Abdelhamid, Y.; Khairy, K.; Abulnoor, B.A.; Ibrahim, S.H.; Khaled, H. The innovation of AI-based software in oral diseases: clinical-histopathological correlation diagnostic accuracy primary study. *BMC Oral Health* **2024**, *24*, 598. Accessed: Mar. 7, 2025, https://doi.org/10.1186/s12903-024-04347-x.