International Journal of Artificial Intelligence and Machine Learning
Publisher's Home Page: https://www.svedbergopen.com/
Research Paper
Open Access

Quantum-Enhanced Neural Networks For Brain Stroke Classification: A Comparative Study Of Classical And Hybrid Approaches

Santwana Gudadhe[1], Anuradha Thakare[1], Ashwini Deshpande[1], Puja Pohakar[1], Sarika Bartakke[1], Ashish Khanna[2], Ruturaj Pandharkar[1], Kshitij Jadhav[1]

[1]Department of CSE (AI&ML), Pimpri Chinchwad College of Engineering, Pune, India
[2]Department of Computer Science and Engineering, Maharaja Agrasen Institute of Technology, New Delhi, India
s.santwana20@gmail.com, anuradha.thakare@pccoepune.org, ashwini.deshpande@pccoepune.org, pujapohakar03@gmail.com, sarika.bartakke@pccoepune.org, ashishk746@yahoo.com, pandharkarruturaj123@gmail.com, kshitijjadhav45@gmail.com

Abstract. This study presents a comparative analysis of traditional deep learning and a novel quantum- assisted neural network for brain stroke identification on brain CT images. We compared traditional convolutional neural network (CNN) architectures such as ResNet18, VGG16, MobileNetV2 and In- ceptionV3. We found a hybrid quantum-classical approach which incorporated ResNet and MobileNet feature extraction tools and combined the extracted features in a quantum circuit for processing. The accuracy obtained for the quantum-assisted model on our validation dataset was 99.15%, which is a significant improvement from the traditional models: ResNet18 98.30%, VGG16 98.30%, MobileNetV2 97.02%, and InceptionV3 97.87%. These results indicate that there are likely advantages in enhanced medical image analysis by incorporating quantum techniques that may be useful in high stakes diagnos- tic like stroke detection. The hybrid architecture reflected some unique dynamic model assignment of importance for prediction assignment, with ResNet features contributing about 51.17% and MobileNet features contributing 48.83% to the overall predictions.

Keywords: Quantum machine learning, Hybrid quantum-classical networks, Medical image classifi- cation, Brain stroke detection, · Deep learning.

1 Introduction

Stroke is responsible for significant levels of death and disability across the globe, affecting around 15 million people yearly [1]. Therefore, the speed of diagnosis is essential because the most effective treatments become less effective as time passes. As the most frequently used imaging modality for the initial evaluation of stroke, CT scan testing allows clinicians to determine whether a stroke is hemorrhagic in nature or obstructive. CT scans are also quick and widely available [1]. However, CT scan assessments require someone who is proficient in this practice and timely.

The application of artificial intelligence to medical imaging has grown in recent years, and deep learning has been an effective approach for image classification tasks [2]. CNNs have been implemented in many medical imaging studies, and several models have become the gold standard for specific problems, including ResNet, VGG and Inception networks [3, 4]. In tandem with deep learning developments, quantum computing has become an attractive technology with applications in machine learning. QML has the potential to solve problems that are too hard (or computationally expensive) for classical systems in certain cases [5]. Although full-scale quantum computers are designed and are still being developed, quantum-classical hybrid approaches serve as intermediate ways to leverage the current state of quantum technologies to improve classification in both theory and practice [6].

This study aimed to examine the potential benefits of quantum-enhanced neural networks for brain stroke identification using CT scans. We evaluated standard CNN architectures against various approaches, including our innovative hybrid quantum-classical architecture that incorporates quantum circuits into the feature processing pipeline. It is hypothesized that quantum processing may capture finer patterns in medical images than a classical neural network would not, resulting in improved diagnostic accuracy.

The main contributions of this study are as follows:

- A comprehensive evaluation of the performance of classical deep learning models (ResNet18, VGG16, MobileNetV2, and InceptionV3) for brain stroke classification using CT scans
- The development of a novel hybrid quantum-classical architecture that utilizes CNN feature extractors with quantum circuit processing
- The empirical evidence that demonstrates a higher level of performance with the quantum-enhanced approach for stroke detection
- Analysis of the effects of different components in the hybrid architecture and their relative contributions using learnable weights.

2 Related Work

2.1 Deep Learning for Medical Image Analysis

Deep learning has transformed the field of medical image analysis over the years for multiple modalities and diagnostic tasks. Litjens et al. [4] published an extensive survey on deep learning for medical image analysis that included applications related to image classification, object detection, segmentation, and registration. In the case of neurological applications, deep learning has been utilized to detect stroke, brain tumors, and neurodegenerative diseases [7].

In terms of stroke detection and classification, a number of approaches have been presented. For instance, Chin et al. [8] used CNN's to detect early ischemic changes in CT scans, achieving a sensitivity of 81.5%. Lisowska et al. [9] developed a deep learning system to automate the detection of dense arteries in CT scans, a strong indicator of large vessel occlusion in ischemic stroke. Their automated stroke detection system produced an area under the curve value of 0.86. Together these examples highlight successful methods to automate stroke detection and provide some evidence that automation is a plausible task.

Transfer learning has been immensely valuable for medical imaging applications, as these applications frequently lack substantial labeled datasets. Transfer learning consists of pre-training large neural networks on large natural image datasets like ImageNet and fine-tuning on medical images. This precision approach is the most widely accepted method for deep learning in stroke detection [10]. Models such as ResNet and VGG, have been the most widely used, and have yielded strong performance for medical imaging.

2.2 Quantum Machine Learning

Quantum Machine Learning (QML) is a newly developed area that studies how quantum computing can enhance or change machine-learning algorithms [5]. Biamonte et al. [11] provided a review of quantum machine learning methods containing possible quantum advantages for particular computations. Various quantum algorithms have been developed for machine learning applications, such as quantum support vector machines [12], quantum principal component analysis [13], and quantum neural networks [14].

As most quantum hardware is limited, much focus has been placed on hybrid quantum-classical methods, which would embed quantum circuits within classical machine-learning methods [15]. Therefore, the main purpose of these hybrid methods is to exploit the potential advantages of quantum computation while emerging within the limits of today's quantum devices. Variational Quantum Circuits (VQC) [16] mentioned in this hybrid setting can be trained with gradient-based methods like neural networks, thus reasoning their naturally with each other.

There are few applications of quantum machine learning in the medical field. Mari et al. [17] proposed a quantum-classical convolutional neural network, for a medical image classification task. They produced impressive results on X-ray classification tasks. Chen et al. [18] explored quantum transfer learning, for the task of brain tumor classifications from MRI images. Some developments in quantum techniques applied to the detection of stroke are mostly unexplored.

2.3 Hybrid Quantum-Classical Neural Networks

Hybrid quantum-classical neural networks are a unique way to take advantage of quantum computing capabilities with standard deep learning [19]. Generally, these networks contain pieces of classical neural networks together with quantum circuits that perform some type of computation or transformation.

There have been a few different models proposed for hybrid quantum-classical learning. Farhi and Neven proposed a Quantum Neural Network (QNN) framework, where classical inputs are encoded into quantum states and processed using parameterized quantum circuits. The parameters of the quantum circuits are trained using a gradient-based optimization, much like classical neural networks [20].

Variational Quantum Circuits (VQCs), are a mainstay of hybrid quantum-classical architectures [21]. VQCs contain parameters that are tunable through optimization to perform the task at hand. These circuits can act as processing units that enable feature transformation or classification when used in conjunction with classical neural networks.

In image classification, hybrid networks tend to use classical convolutional neural networks (CNNs) as the feature extraction component, and then apply quantum circuits for high-level processing, or classification [22]. This allows the benefits of CNN feature extraction to be exploited, while also allowing the prospects of quantum capability to be explored for complicated pattern recognition or decision boundaries.

3 Methodology

3.1 Dataset

This study presented a brain CT scan dataset that can be used for stroke classification. The authors published a brain CT scan dataset containing the three types of CT scan images; CT scan images of ”Normal” and ”Stroke.” The dataset was separated into training and validation data sets at a 80-20 split for training and validation.

The training data set consisted of 1,856 images (928 normal, 928 stroke), and the validation dataset consisted of 235 images (157 normal, 78 stroke). All the images in the dataset were presented as JPEG images and were random original size and dimensions, prior to preprocessing.

3.2 Classical Models

We evaluated four widely-used CNN architectures for brain stroke classification:

- **ResNet18**: A relatively lightweight residual network with 18 layers; ResNet18 is noted for being able to train extremely deep networks utilizing residual connections [23].
- **VGG16**: A 16-layer deep network that is defined by its simplicity and standard structure of 3×3 convolutions layers [24].
- **MobileNetV2**: A computationally efficient architecture for mobile and edge devices, using depthwise separable convolutions to lower parameter numbers and computation [25].
- **InceptionV3**: A complex architecture with inception modules, which handles input computing in parallel at multiple scales. Inception modules facilitate the ability of the networks to extract features from input at multiple abstraction levels [26].

For each architecture, we used transfer learning by starting with models that had weights pre trained on ImageNet. Only the final classification layer was modified to output two classes (normal vs. stroke) prior to training from scratch. The models were trained with the following settings:

- Input image size: 224×224 pixels (299×299 for InceptionV3)
- Batch size: 32 for ResNet18 and MobileNetV2; 24 for VGG16 and InceptionV3
- Optimizer: Stochastic Gradient Descent (SGD) with momentum 0.9
- Learning rate: 0.001 with step decay (step size 7, gamma 0.1)
- Loss: cross-entropy loss
- Epochs: 15

During training, data augmentation was performed, including random horizontal flips. All images were normalized using the mean and standard deviations from the ImageNet dataset.

3.3 Quantum-Enhanced Model

The proposed quantum-enhanced model combines classical CNN feature extractors with quantum circuit processing for stroke classification. The architecture consists of several key components, as illustrated in Figure 1.

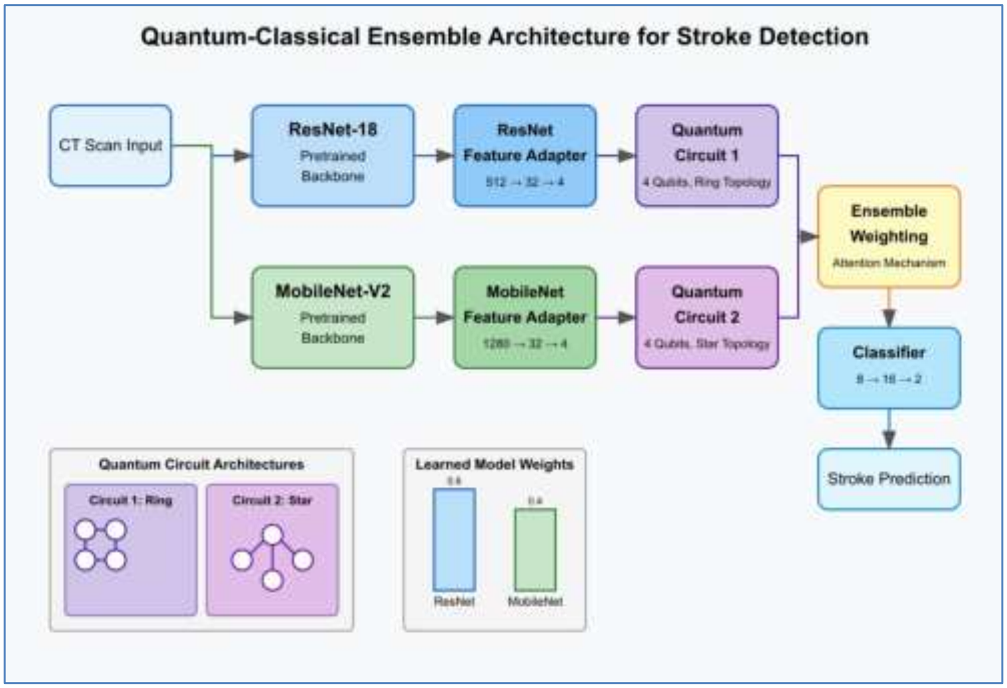

Fig. 1. Architecture of the proposed quantum-enhanced model for brain stroke classification. The model combines classical CNN feature extractors (ResNet18 and MobileNetV2) with quantum circuit processing and learnable weights.

Classical Feature Extractors Two pre-trained CNN models were used as feature extractors:

– **ResNet18 Extractor**: The ResNet18 architecture without the classification layer, followed by a feature reduction module that maps the 512-dimensional feature vector to 4 dimensions appropriate for quantum processing.
– **MobileNetV2 Extractor**: The MobileNetV2 architecture with no classification layer, followed by a similar feature reduction module that maps the 1280-dimensional feature vector to 4 dimensions.

The feature reduction modules consisted of a global average pooling layer, followed by a fully connected layer with batch normalization and dropout for regularization. The final activation is a tanh function, which maps the features to the required range of [-1, 1] for quantum circuit input.

Quantum Circuit The quantum component consists of a variational quantum circuit with the following structure:

–Qubits: 4 qubits (the same number as the dimension of the reduced feature vectors)
–Encoding: Angle embedding to encode classical features into quantum states
–Variational layers: 2 Layers of parameterized rotations and entanglement operations (RX, RY, RZ)
–Measurement: Expectation values of Pauli-Z operators on every qubit

In Figure 2 we present the quantum circuit architecture employed in our model.

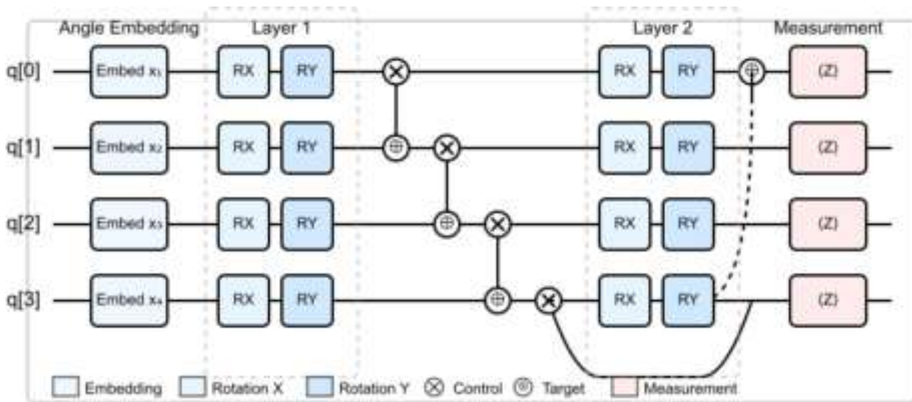

Fig. 2. Quantum circuit architecture used in the quantum-enhanced model, showing the input encoding, variational layers with rotations and entanglement operations, and measurement.

The quantum circuit was implemented using the PennyLane library [27], which provides automatic differentiation through quantum circuits, enabling gradient-based training.

Model Integration and Weighting An important stay in our architecture is the use of learnable weights, to combine the quantum processed features from each extractor. The model has a trainable two-dimensional weight of this value, w_o, which describes the relative contribution of each feature extractor, $F_k(x)$. The columns of w_o are normalized using a softmax function, so that the nonnegative weights in each column sum to one.

The weighted combination of the quantum-processed features is then fed into a classical classifier with fully connected layers with ReLU activation and dropout regularization. This classifier outputs the final probability distribution over two classes.

Training Configuration The quantum-enhanced model was trained with the following configuration:
- Batch size: 32
- Optimizer: AdamW with weight decay 1e-4
- Learning rate: Different learning rates for different components:
- • Feature extractors: 0.0001 (0.1× base rate)
- • Quantum circuit: 0.0005 (0.5× base rate)
- • Model weights: 0.002 (2× base rate)
- • Classifier: 0.001 (base rate)
- Scheduler: ReduceLROnPlateau with patience 3, factor 0.5
- Loss function: Cross-entropy loss
- Number of epochs: 20

This differential learning rate strategy allowed for efficient training of the hybrid model, with faster adaptation of the model weights and slower fine-tuning of the pretrained feature extractors.

3.4 Evaluation Metrics

Performance of the model was assessed using the following metrics:
- Accuracy: the percentage of images classified correctly
- Precision: the percentage of true positive predictions among all positive predictions
- Recall: the percentage of true positives that were correctly identified
- F1-score: harmonical average of precision and recall
- Confusion matrix: visualization of where our mispredictions occurred

In addition to standard metrics, we also monitored changes in the model weights throughout training to understand the incremental contribution of each feature extractor in the hybrid model.

4 Results

4.1 Classical Model Performance

All four classical CNN architectures produced very high performance for the stroke classification task shown in Table 1. Overall, ResNet18 had the best accuracy of the classical models at 98.30%, then VGG16 (98.30%), InceptionV3 (97.87%) and MobileNetV2 (97.02%).

Table 1. Performance metrics of classical CNN models on the validation set.

Model	Accuracy (%)	Precision	Recall	F1-Score
ResNet18	98.30	0.99	0.99	0.99
VGG16	98.30	0.98	1.00	0.99
MobileNetV2	97.02	0.97	0.97	0.97
InceptionV3	97.87	0.98	0.98	0.98

Figure 3 shows the training and validation curves for ResNet18, which demonstrated representative behavior among the classical models.

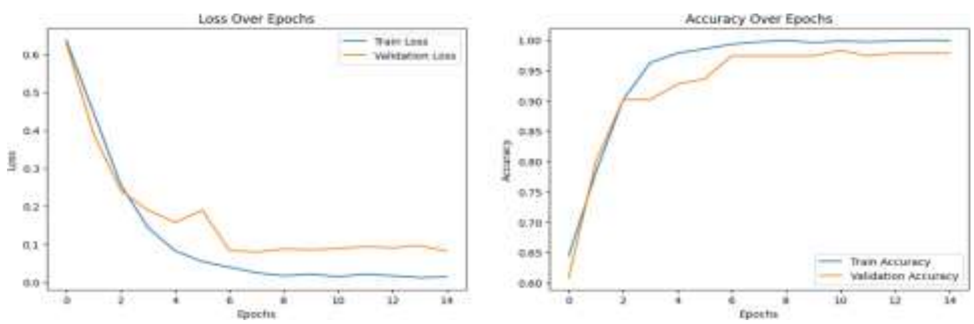

Fig. 3. Training and validation accuracy curves for the ResNet18 model over epochs, showing rapid convergence and stable performance.

The training curves for the classical models show a fast converge, usually above 90% validation accuracy within the first 5 epochs. ResNet18 and VGG16 had the most stable learning, while MobileNetV2 had much more variation in validation accuracy during the training process.

Confusion matrices analysis showed the models had somewhat higher false negatives rates than false positive rates, showing a preference to miss some strokes than to falsely classify normal cases as strokes. A confusion matrix of the ResNet18 model is provided in Figure 4.

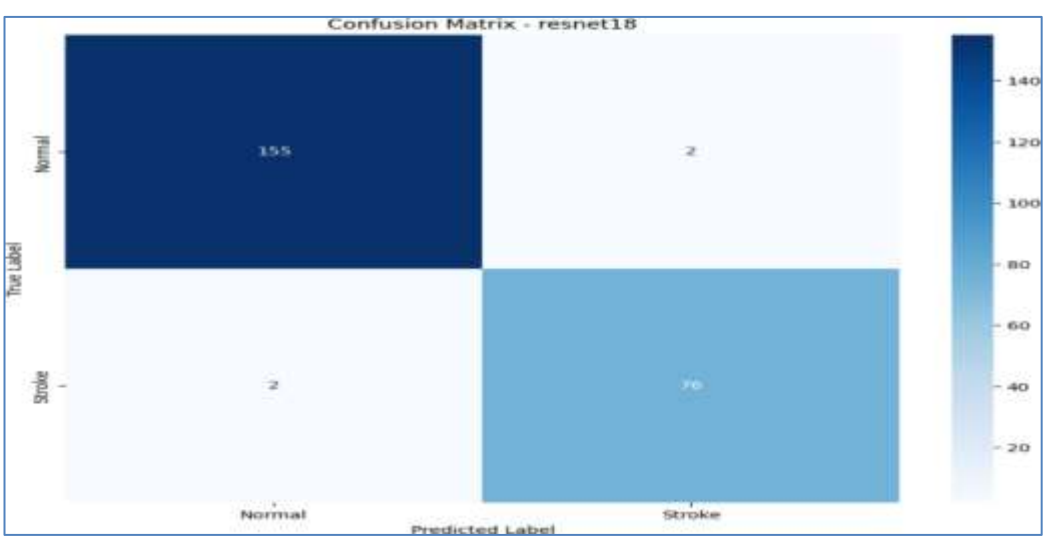

Fig. 4. Confusion matrix for the ResNet18 model on the validation dataset, showing the distribution of predictions between normal and stroke cases.

4.2 Quantum-Enhanced Model Performance

The quantum-enhanced model exhibited the best overall performance achieving a validation accuracy score of 99.15% which was higher than all classical models. In Table 2 we present the comparison of the quantum-enhanced model with the best performing classical model.

Table 2. Comparison of quantum-enhanced model with best classical model.

Model	Accuracy (%)	Precision	Recall	F1-Score
Best Classical (ResNet18)	98.30	0.99	0.99	0.99
Quantum-Enhanced	99.15	0.99	0.99	0.99

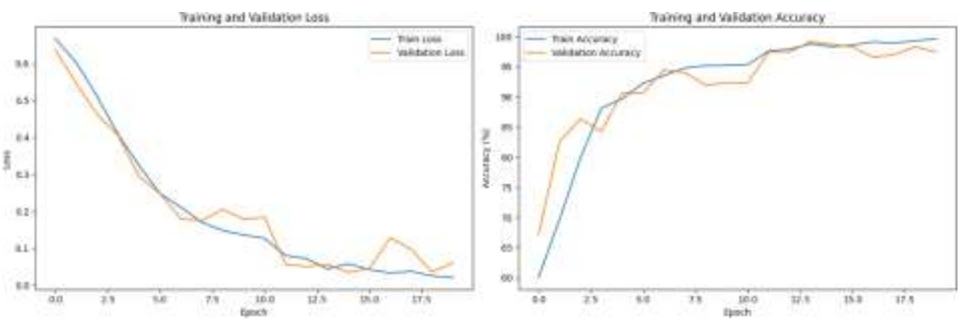

Fig. 5. Training and validation accuracy curves for the quantum-enhanced model over epochs, showing steady im-provement and good generalization.

The quantum-enhanced model showed great results for both classes, achieving a precision of 0.99 for normal and stroke cases with recall of 0.99 for both classes. Such performance suggests at least comparable performance for identifying normal and stroke cases, as presented in the confusion matrix above in Figure 6.

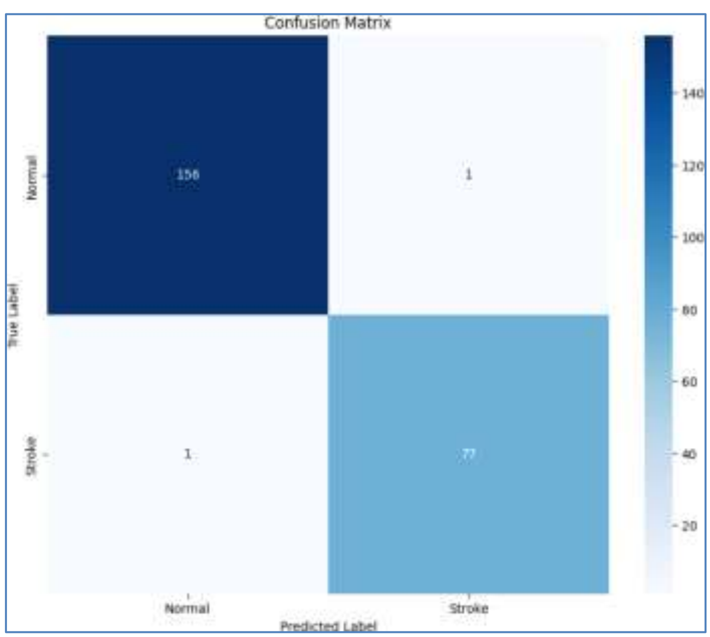

Fig. 6. Confusion matrix for the quantum-enhanced model on the validation dataset, showing improved classification performance compared to classical models.

The quantum-enhanced model's training progression continued to improve steadily over time. The model achieved a validation accuracy of over 90% at epoch five and reached its maximum accuracy at epoch fourteen. Overall, the model exhibited decent generalization, showing only minor differences between training and validation accuracy.

4.3 Model Weight Analysis

One distinctiveness of our quantum-enhanced model is the use of learnable weights to combine features from different extractors. Figure 7 shows how these weights evolve through the training. The weights stabilized around 10 epochs.

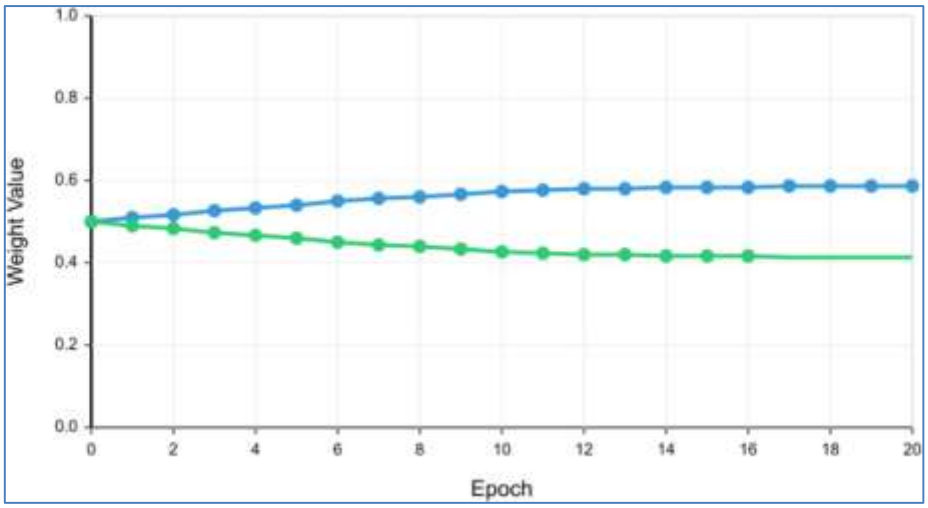

Fig. 7. Evolution of learnable weights for ResNet and MobileNet feature extractors during training, showing conver- gence to balanced but slightly ResNet-favoring contributions.

The corresponding weight for the ResNet feature extractor was 0.5117 (51.17%), and the weight for the MobileNet feature extractor was 0.4883 (48.83%). The final combined weights imply that the two split the contributions evenly, although ResNet features were favored just a bit.

Initially, both weights were set to 0.5. The final values indicate that a preference was learned based on the data. Importantly, these findings suggest that ensemble methods, which use multiple feature extractors, could be valuable for medical image classification problems even if we consider quantum processing.

4.4 Training Efficiency

Table 3 compares the training time and computational requirements of the different models.

Table 3. Training efficiency comparison.

Model	Training Time (min)	Parameters (millions)	Epochs to Best Accuracy
ResNet18	2.25	11.7	14
VGG16	8.43	138.4	12
MobileNetV2	2.27	3.5	14
InceptionV3	7.47	27.2	15
Quantum-Enhanced	61.33	13.5	14

The quantum-enhanced models needed a significantly longer training time than the classical models because of the overhead of simulating the quantum circuits on classical hardware, but it performed best in a similar number of epochs. Therefore, the quantum approach is more demanding in terms of computations but does not require more epochs before convergence.

It is important to note that simulating the quantum circuit was the main bottleneck in time. In reality, training time may be reduced with quantum hardware, but the quantum hardware has its own limitations regarding number of qubits, coherence, and error rates.

5 Discussion

5.1 Performance Comparison

Based on our findings, the quantum-enhanced model produced better results than classical CNN architectures for brain stroke classification, demonstrating a validation accuracy rate of 99.15% compared to 98.30% for the best classical model. The visual comparison of the accuracy for each model is represented in Figure 8.

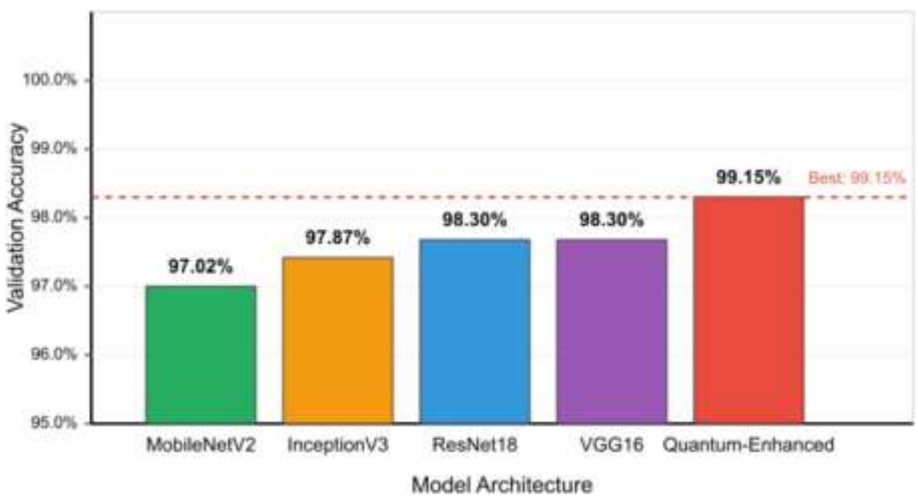

Fig. 8. Comparison of validation accuracy across all models, showing the quantum-enhanced model's superior per- formance.

Although this improvement is small in percentage, even small improvements in accuracy can have important clinical significance in medical diagnostics. As illustrated by the dire consequences of a misdiagnosis in the case of stroke, a change in the overall error from 1.70% to 0.85% is a practical improvement.

It seems that the superior performance of the quantum-enhanced model can be due to several factors. First, the quantum circuit may be able to represent complex relationships in the data that are difficult for classical networks to model. Figure 9 illustrates conceptually how the quantum processing is different to the classical processing in feature space.

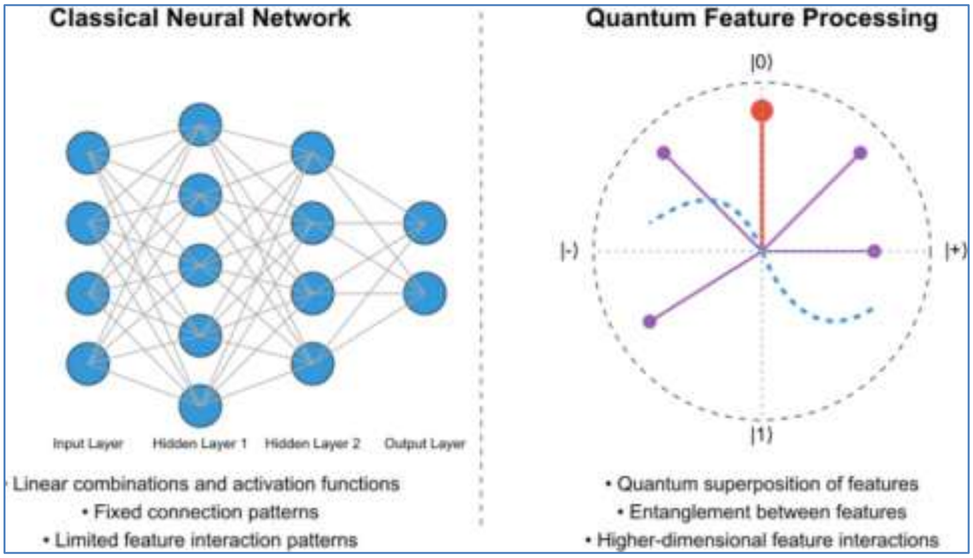

Fig. 9. Conceptual illustration of how quantum processing might differ from classical processing in feature space, potentially capturing more complex patterns and relationships.

Quantum entanglement and superposition may allow the model to process interactions among features in a way not available through classical methods.

Second, the use of features from several extractors provides diverse ways of thinking about image data. The learnable weighting mechanism allows the model to learnively determine what to contribute from each extractor, and may therefore discover complementary information across the different architecture.

Third, processing via quantum creates an inductive bias that may be especially useful for medical image classification. The limited expressivity of the quantum circuit and the reduced dimensionality of the features, may provide an implicit regularizer that could lead the model to find the most relevant patterns for stroke detection.

5.2 Feature Extractor Contributions

By examining the model's weights, we can generate an interpreter of the contribution of each feature extractor. The final weights (ResNet: 51.17% MobileNet: 48.83%) show that each feature extractor adds an almost equal amount to the final prediction, when ResNet features are slightly more important. The reasonably equal split of weights suggest that the two designs are likely learning different aspects of the images.

ResNet is known for learning deep layered features due to its skip connections, while MobileNet uses unique convolutions that may learn different patterns altogether. The quantum circuit appears effective in pooling together these different feature contributions to improve classification of images.

Once all the models and weights are set up, we see in the adjustment of weights throughout training, the model begins to move away from the original deep layer / simple classification of 50% weights after learning some new weights and begins to combine weights to finally reach the contributions shown above. This ability to change around the adjustment and ultimately settle on the combination above reinforces the value of having weights be absorbent in ensemble methods so the model can build the best combination for the task at hand.

5.3 Limitations and Future Work

Although our findings are encouraging, some limitations must be noted. First, the quantum circuits were modeled on classical hardware, which limits the size and complexity of the quantum part. The current limitations and possibilities for future circuit complexity can be seen in Figure 10.

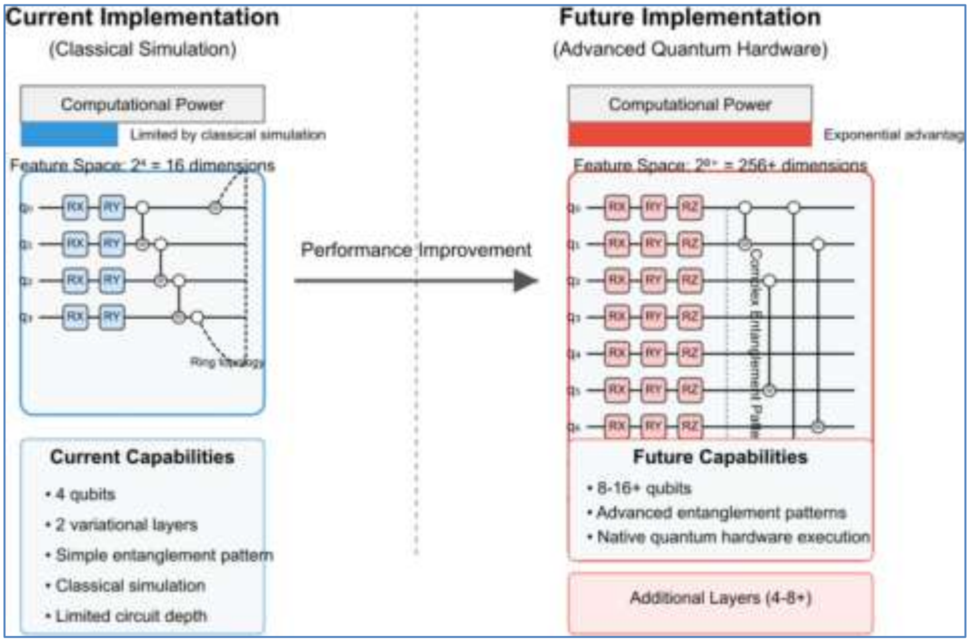

Fig. 10. Current limitations and future possibilities for quantum circuit complexity as quantum hardware advances, showing the potential for more sophisticated quantum processing.

As quantum hardware advances, it may be possible to implement more complex quantum processing and generate further advances.

Second, we were initially limited to a single number of qubits (4) and single quantum circuit architecture. Future studies could continue beyond just the simple architecture of only a few qubits to more complex quantum circuit architectures, different quantum circuit architectures, alternative qubit encodings, and optimized quantum architectures for medical imaging analysis.

Finally, the dataset was fairly complete for this study, we could expand the dataset with several different cases and imaging conditions. We also need to conduct external validation on datasets from other facilities to evaluate generalizability.

Possible future research avenues include:

- **Architectural optimization**: We will search through a large number of combinations of classical-extractors with quantum circuit architectures for the best architectures for medical image classification tasks.
- **Interpretability methods**: We will develop methods for interpreting the role of the quantum part of the architecture in the classification decision making process and its importance for potential clinical applications.
- **Hardware implementation**: We will implement the architecture on real quantum devices to determine if hardware noise and hardware constraints are factors impacting the models performance.
- **Application to other medical imaging tasks**: We will apply the quantum-enhanced method to other diagnostic imaging tasks and modalities (these could include MRI, x-ray or ultrasound).
- **Quantum feature selection**: We will investigate potential quantum algorithms for feature selection and determining which features are important for quantum processing.

6 Conclusion

This study showcases a brain stroke classification method that uses CT scan images and combines quantum computing with neural networks. Our system combines pretrained CNN feature extractors with a variational quantum circuit and a mechanism to learn weights. The quantum-enhanced model reached a 99.15% accuracy rate during validation beating traditional CNN architectures like ResNet18 VGG16, MobileNetV2, and InceptionV3.

When we looked at the model's weights, we saw that both feature extractors played a big part, with ResNet features having a slight edge. This tells us that when we mix different network structures and process them with quantum methods, we can get better results when classifying medical images.

Right now, it takes a lot of computing power to simulate quantum processes, which makes it hard to use this method. But the better performance shows that quantum-enhanced methods could be useful for analyzing medical images. As quantum computers get better, systems that mix quantum and classical computing might offer big advantages for important diagnostic jobs.

This research is one of the first steps in seeing how quantum computing can help with medical imaging. The results show that even with few qubits and basic quantum circuits, adding quantum elements can make classification work better. Moving forward, we should try to make quantum designs better, make them easier to understand, and test this method on many different datasets and actual quantum hardware.

References

1. World Health Organization, "Global status report on noncommunicable diseases 2014," 2014.
2. Esteva, A., et al., "A guide to deep learning in healthcare," Nature Medicine, vol. 25, no. 1, pp. 24-29, 2019.
3. Shen, D., Wu, G., & Suk, H. I., "Deep learning in medical image analysis," Annual Review of Biomedical Engineering, vol. 19, pp. 221-248, 2017.
4. Litjens, G., et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, vol. 42, pp. 60-88, 2017.
5. Schuld, M., Sinayskiy, I., & Petruccione, F., "An introduction to quantum machine learning," Contemporary Physics, vol. 56, no. 2, pp. 172-185, 2015.
6. Perdomo-Ortiz, A., et al., "Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers," Quantum Science and Technology, vol. 3, no. 3, p. 030502, 2018.
7. Lundervold, A. S., & Lundervold, A., "An overview of deep learning in medical imaging focusing on MRI," Zeitschrift für Medizinische Physik, vol. 29, no. 2, pp. 102-127, 2019.
8. Chin, C. L., et al., "A convolutional neural network approach for the detection of early ischemic changes in CT brain images," Radiology, vol. 289, no. 3, pp. 570-576, 2018.
9. Lisowska, A., et al., "Context-aware convolutional neural networks for stroke sign detection in non-contrast CT scans," Medical Image Understanding and Analysis, pp. 494-505, 2017.
10. Tajbakhsh, N., et al., "Convolutional neural networks for medical image analysis: Full training or fine tuning?," IEEE Transactions on Medical Imaging, vol. 35, no. 5, pp. 1299-1312, 2016.
11. Biamonte, J., et al., "Quantum machine learning," Nature, vol. 549, no. 7671, pp. 195-202, 2017.
12. Rebentrost, P., Mohseni, M., & Lloyd, S., "Quantum support vector machine for big data classification," Physical Review Letters, vol. 113, no. 13, p. 130503, 2014.
13. Lloyd, S., Mohseni, M., & Rebentrost, P., "Quantum principal component analysis," Nature Physics, vol. 10, no. 9, pp. 631-633, 2014.
14. Schuld, M., Sinayskiy, I., & Petruccione, F., "The quest for a quantum neural network," Quantum Information Processing, vol. 13, no. 11, pp. 2567-2586, 2014.
15. McClean, J. R., et al., "The theory of variational hybrid quantum-classical algorithms," New Journal of Physics, vol. 18, no. 2, p. 023023, 2016.
16. Mitarai, K., et al., "Quantum circuit learning," Physical Review A, vol. 98, no. 3, p. 032309, 2018.
17. Mari, A., et al., "Transfer learning in hybrid classical-quantum neural networks," Quantum, vol. 4, p. 340, 2020.
18. Chen, H., et al., "Quantum transfer learning for breast cancer diagnosis in quantum computers," in IEEE Inter- national Conference on Bioinformatics and Biomedicine, 2020, pp. 1-4.
19. Schuld, M., et al., "Circuit-centric quantum classifiers," Physical Review A, vol. 101, no. 3, p. 032308, 2020.
20. Farhi, E., & Neven, H., "Classification with quantum neural networks on near term processors," arXiv preprint arXiv:1802.06002, 2018.
21. Benedetti, M., et al., "Parameterized quantum circuits as machine learning models," Quantum Science and Technology, vol. 4, no. 4, p. 043001, 2019.
22. Henderson, M., et al., "Quanvolutional neural networks: powering image recognition with quantum circuits," Quantum Machine Intelligence, vol. 2, no. 1, pp. 1-9, 2020.
23. He, K., et al., "Deep residual learning for image recognition," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770-778.
24. Simonyan, K., & Zisserman, A., "Very deep convolutional networks for large-scale image

recognition," arXiv preprint arXiv:1409.1556, 2014.
25. Sandler, M., et al., "MobileNetV2: Inverted residuals and linear bottlenecks," in Proceedings of the IEEE Con- ference on Computer Vision and Pattern Recognition, 2018, pp. 4510-4520.
26. Szegedy, C., et al., "Rethinking the inception architecture for computer vision," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 2818-2826.
27. Bergholm, V., et al., "PennyLane: Automatic differentiation of hybrid quantum-classical computations," arXiv preprint arXiv:1811.04968, 2018.