



**IJITCE**

**ISSN 2347- 3657**

# International Journal of Information Technology & Computer Engineering

[www.ijitce.com](http://www.ijitce.com)



Email : [ijitce.editor@gmail.com](mailto:ijitce.editor@gmail.com) or [editor@ijitce.com](mailto:editor@ijitce.com)

# Quantum-Enhanced Approach for Brain Tumor Classification

Dr. P. Dhana Lakshmi, D. Vamsi, D. Uma Siva Shankar , M. L. R. Sai Harshit

Department of Electronics and Communication Engineering  
Acharya Nagarjuna University College of Engineering and Technology  
Nagarjuna Nagar, Guntur, Andhra Pradesh, India

**Abstract**—This paper investigates a quantum-enhanced approach for brain tumor classification using MRI data. We focus on establishing robust classical deep learning baselines and defining the core components for integrating Quantum Convolutional Neural Networks (QCNNs). Classical models, specifically VGG-16, MobileNetV1, and Xception, are employed via transfer learning on a prepared brain tumor dataset to benchmark performance. Standardized preprocessing from .mat files, data augmentation, and training procedures are detailed. Concurrently, we define the architecture for a potential hybrid QCNN system using PennyLane, featuring a 4-qubit variational quantum circuit designed as a Keras Layer for integration with classical feature extractors. The objective is to lay the groundwork for exploring how quantum-inspired processing, combined with powerful classical feature extraction, might enhance pattern recognition and classification accuracy for this critical medical imaging task, particularly in complex cases. This work provides the necessary data handling protocols, baseline results, and quantum component specifications for future implementation and evaluation of the complete hybrid system.

**Index Terms**—Brain Tumor Classification, MRI, Deep Learning, Transfer Learning, VGG-16, MobileNet, Xception, ResNet, Quantum Convolutional Neural Networks (QCNN), Hybrid Quantum-Classical Models, Quantum Computing, Medical Image Analysis, PennyLane.

## I. INTRODUCTION

Brain tumor classification is a critical task in medical diagnostics, with Magnetic Resonance Imaging (MRI) being a primary non-invasive modality [12]. Accurate classification of tumor types (e.g., meningioma, glioma, pituitary) is essential for treatment planning and prognosis [13]. While human interpretation is standard, it can be subjective. Deep Learning (DL), especially Convolutional Neural Networks (CNNs), has shown remarkable success in medical image analysis [14], [15], often leveraging transfer learning with architectures like VGG [19], MobileNet [20], ResNet [21], and Xception [22] to overcome data scarcity [16].

Despite advances, classical DL models can face challenges, particularly with subtle image features or ambiguous cases. Concurrently, quantum computing offers a new paradigm for computation, with Quantum Machine Learning (QML) [2], [3] exploring its potential benefits. Quantum Convolutional Neural Networks (QCNNs) [1], [4] adapt CNN concepts to the quantum domain, potentially leveraging quantum phenomena like superposition and entanglement to capture complex data correlations [8].

Due to the limitations of current Noisy Intermediate-Scale Quantum (NISQ) devices [9], purely quantum approaches for high-resolution imaging are impractical. Hybrid quantum-classical models [5], [6] offer a promising near-term direction. These models typically use classical CNNs for initial feature extraction and dimensionality reduction, followed by a smaller Parameterized Quantum Circuit (PQC) or Variational Quantum Circuit (VQC) for further processing or classification [7]. Frameworks like PennyLane [17] facilitate building and training such models by bridging classical frameworks (like TensorFlow [18]) and quantum backends.

This paper focuses on the essential preparatory steps towards developing such a hybrid QCNN for brain tumor classification. We systematically preprocess MRI data, establish strong classical performance baselines using standard transfer learning techniques, and meticulously define the components of a 4-qubit QCNN layer intended for integration. This foundational work establishes the necessary benchmarks and components required for the subsequent development and rigorous evaluation of the complete hybrid system against purely classical approaches.

## II. RELATED WORK

The field of automated brain tumor classification spans classical machine learning, deep learning, and emerging quantum approaches.

### A. Classical Methods

Early approaches often relied on handcrafted features extracted from MRI images, such as texture, shape, or intensity statistics, followed by traditional classifiers like Support Vector Machines (SVMs) or decision trees [12]. While interpretable, these methods often require significant domain expertise and may lack robustness to variations in image acquisition or tumor appearance.

### B. Deep Learning Approaches

The advent of deep learning, particularly CNNs, marked a significant shift [14], [15]. CNNs learn hierarchical features directly from pixel data, reducing the need for manual feature engineering. Transfer learning [13], [16], using models pre-trained on large datasets like ImageNet (e.g., VGG [19], MobileNet [20], ResNet [21], Xception [22]), has become a standard technique, especially when medical datasets are limited. These models learn general visual patterns that can be

adapted effectively to medical imaging tasks. Variants like 3D CNNs or combinations with recurrent networks (LSTMs) have also been explored [16].

### C. Quantum Machine Learning Approaches

QML is a nascent field exploring quantum algorithms for ML tasks [2], [3]. QCNNs [1], [4] are quantum analogues of classical CNNs, using quantum circuits for feature extraction. Variational Quantum Circuits (VQCs) [6] are often used, where circuit parameters are optimized classically. Hybrid models [5], [7] combine classical deep learning for initial processing with VQCs for potentially enhanced classification, representing a practical approach for NISQ-era devices [9]. Frameworks like PennyLane [17] enable the development and automatic differentiation of these hybrid computations. Several reviews discuss the potential and challenges of QML in medical imaging [25], [26]. Our work aligns with this hybrid paradigm, focusing on defining and preparing for such a system.

## III. DATASET AND PREPROCESSING

### A. Dataset Description

This study utilizes a publicly available brain tumor dataset [12], containing T1-weighted contrast-enhanced MRI images across three tumor types: meningioma (label 1), glioma (label 2), and pituitary tumor (label 3). The data is originally provided in the MATLAB .mat format. For the experiments reported here, a subset of 250 .mat files was initially processed. Each file includes the image data (accessed via `cjdata['image']`) and the corresponding label (`cjdata['label']`). The full dataset comprises significantly more images (e.g., 3064 samples).

### B. Preprocessing Pipeline

A standardized preprocessing pipeline was implemented to prepare the data for deep learning models:

- 1) *File Reading & Format Conversion*: The `h5py` library read image arrays and labels from .mat files. Original grayscale images were converted to 3-channel RGB by duplicating the channel, ensuring compatibility with ImageNet pre-trained models. Images were saved losslessly as .jpg using `matplotlib.pyplot.imshow` without axes or padding.
- 2) *Directory Organization*: Images were sorted into class-specific directories ('1', '2', '3') based on labels for compatibility with Keras directory iterators.
- 3) *Dataset Splitting*: The `splitfolders` library partitioned the .jpg images into training (80%) and validation (20%) sets, stratified by class. A random seed (42) ensured reproducibility. This yielded 198 training and 52 validation images from the initial 250 files.
- 4) *Image Augmentation (Training Set Only)*: TensorFlow's `ImageDataGenerator` applied on-the-fly augmentation to the training data to enhance robustness and mitigate overfitting: random rotations (up to 45°), random horizontal flips, and random vertical flips.

- 5) *Resizing and Normalization*: All images were resized to 224x224 pixels. Pixel values were normalized to the [0, 1] range by dividing by 255.0.
- 6) *Data Generators*: Keras `ImageDataGenerator` instances with `flow_from_directory` were used to load images in batches (size 32) from the respective train and val directories. `class_mode='categorical'` provided one-hot encoded labels for training with categorical crossentropy loss.

## IV. PROPOSED HYBRID QCNN ARCHITECTURE (CONCEPTUAL)

We propose a hybrid quantum-classical architecture designed to leverage the strengths of both paradigms for brain tumor classification. The core idea is to use established classical CNNs for robust feature extraction from high-dimensional MRI data and then employ a compact quantum circuit to potentially capture complex correlations within these extracted features.

### A. Classical Feature Extractor Component

The classical part utilizes pre-trained CNNs as feature extractors.

- 1) *Base Model*: Architectures like VGG16 [19], MobileNetV1 [20], or Xception [22] are loaded with ImageNet weights, excluding the final classification layer (`include_top=False`). Input shape is set to (224, 224, 3).
- 2) *Feature Extraction Mode*: The layers of the base model are initially frozen (`trainable = False`) to act as fixed feature extractors. Optionally, later layers could be unfrozen for fine-tuning during hybrid training.
- 3) *Interface Head*: A custom sequence of Dense layers is added on top of the base model's output (after global pooling, e.g., `GlobalMaxPooling2D`). This head performs further classical processing and dimensionality reduction, culminating in a `Dense(4)` layer that produces a 4-dimensional feature vector. This vector serves as the input interface to the quantum component.

### B. Quantum Circuit Component (QCNN Layer)

The quantum part consists of a variational quantum circuit implemented using PennyLane [17].

- 1) *Quantum Device*: A simulator, such as PennyLane's `lightning.qubit` or `default.qubit`, is instantiated with 4 qubits.
- 2) *QNode Definition*: A Python function decorated with `@qml.qnode` defines the quantum computation:
  - **Encoding**: The 4 classical features from the `Dense(4)` layer are encoded into the quantum state using `qml.AngleEmbedding`. This maps each feature  $x_i$  to a rotation angle  $\phi_i$  applied to a corresponding qubit, e.g., via  $R_Y(\phi_i)$  gates.
  - **Ansatz (Variational Circuit)**: A sequence of parameterized quantum gates forms the trainable part of the circuit. We use `qml.BasicEntanglerLayers`






Fig. 1. Conceptual Diagram of the Proposed Hybrid Quantum-Classical Architecture.

(e.g., 3 layers), which apply parameterized single-qubit rotations (RX, RY, RZ) to each qubit, followed by entangling CNOT gates (e.g., in a ring topology). The rotation parameters are the trainable weights of the quantum layer.

- **Measurement:** Expectation values of the Pauli-Z operator,  $\langle \hat{Z}_i \rangle$ , are measured for each of the 4 qubits. This yields a 4-dimensional classical output vector representing the quantum-processed features.

3) **Keras Layer Integration:** The QNode is wrapped into a `qml.qnn.KerasLayer` (or using newer PennyLane integration methods). This allows the quantum circuit to be seamlessly inserted into a Keras Sequential model, specifying the trainable weight shapes (e.g., `weight_shapes={"weights": (num_layers, num_qubits)}`) and the output dimension (4).

### C. Overall Hybrid Model Structure

The complete conceptual model connects these components: Input Image (224x224x3) → Classical Base Model (Frozen) → Global Pooling → Custom Dense Head (including Dense(4)) → KerasLayer(QNode) → Final Dense(3, activation='softmax') layer for 3-class classification output. Fig. 1 provides a conceptual illustration.

### D. Training Strategy

The entire hybrid model is intended to be trained end-to-end. Gradients are calculated via backpropagation, seamlessly flowing through both classical Keras layers and the quantum Keras layer using PennyLane's automatic differentiation capabilities (e.g., parameter-shift rule or adjoint method). A standard optimizer like Adam [23] would update both the classical parameters in the head (and potentially unfrozen base layers) and the quantum parameters in the ansatz.

## V. IMPLEMENTATION DETAILS & EXPERIMENTAL SETUP

### A. Platform and Setup

Experiments were conducted using Python 3.x in an environment like Google Colab, leveraging TensorFlow 2.x [18]

and PennyLane [17]. GPU acceleration (NVIDIA T4) was utilized for training the classical baseline models. Key libraries included h5py, splitfolders, numpy, matplotlib, and scikit-learn (for potential future evaluation metrics).

### B. Dataset and Configuration

The experiments reported here used the preprocessed subset of 250 images, split into 198 training and 52 validation samples (details in Section III). Images were 224x224x3, normalized to [0, 1]. A batch size of 32 was used for training.

### C. Classical Baseline Training

The classical models (VGG16, MobileNetV1, Xception) were configured as described in Section IV-A.

- **Optimizer:** Adam [23] with a learning rate of 0.001.
- **Loss Function:** Categorical Crossentropy.
- **Metrics:** Accuracy.
- **Epochs:** MobileNetV1 was trained for 60 epochs; VGG16 and Xception were trained for 50 epochs. Base models were kept frozen.

### D. Quantum Component Setup

The quantum layer defined in Section IV-B used:

- **Device:** 4 qubit simulator (lightning.qubit preferred).
- **Ansatz:** 3 layers of `qml.BasicEntanglerLayers`.
- **Gradient Method:** Adjoint differentiation specified within the QNode.
- **Interface:** `qml.qnn.KerasLayer` (Note: Future work may require updating to newer PennyLane-TF integration methods due to deprecation warnings).

## VI. EVALUATION PROTOCOL (FOR BASELINES AND FUTURE HYBRID)

To assess the performance of the classical baselines and provide a framework for evaluating the future hybrid model, we employ the following standard metrics:

- **Overall Accuracy:** The primary metric, representing the percentage of correctly classified samples over the entire validation/test set.
- **Loss and Accuracy Curves:** Plots of training and validation loss/accuracy over epochs to monitor learning dynamics, convergence, and potential overfitting.
- **Per-Class Performance Metrics:** Calculation of Precision, Recall, and F1-Score for each tumor class to understand class-specific performance, particularly important in imbalanced datasets (though stratification was used here).
- **Confusion Matrix:** Visualization of classification results showing correct predictions and specific misclassification patterns between classes.
- **(Future) ROC/PR Analysis:** For more detailed discriminative capability assessment, One-vs-Rest Receiver Operating Characteristic (ROC) curves with Area Under Curve (AUC) and Precision-Recall (PR) curves with Average Precision (AP) can be computed.




Fig. 2. Training and Validation Loss for VGG16 Baseline Model (50 Epochs).

- **(Future) Feature Visualization:** Techniques like t-SNE or PCA could be applied to visualize the features extracted by the classical layers just before the quantum layer, or the features output by the quantum layer, to gain qualitative insights into the representations learned.

These metrics provide a comprehensive view of model performance beyond simple accuracy. The results presented in this paper focus on the loss/accuracy curves and overall validation accuracy for the classical baselines.

## VII. RESULTS AND ANALYSIS (CLASSICAL BASELINES)

### A. Classical Model Performance

The training process yielded classical baseline models with strong performance on the validation set derived from the initial 250 images.

Fig. 2 shows the training and validation loss for the VGG16-based model. The curves decrease steadily and track closely, indicating effective learning and good generalization over 50 epochs on this subset.

The accuracy curves for the MobileNetV1-based model (Fig. 3) show rapid convergence. Training accuracy nears 100%, while validation accuracy jumps to approx. 96.15% early on and then plateaus completely for the remaining 60 epochs. This suggests the model quickly mastered the small validation set, possibly indicating limitations in its diversity or potential early overfitting relative to this specific set.

The Xception-based model's loss curves (Fig. 4), trained for 50 epochs, demonstrate very smooth convergence, with training and validation losses remaining tightly coupled, suggesting excellent generalization for this architecture on this data subset.

Final validation accuracies were: MobileNetV1 **96.15%** (at plateau), VGG16 **96.15%** (at plateau), and Xception **98.08%**

### B. Example Prediction Analysis

Fig. 5 illustrates a prediction by the VGG16 baseline on an image from class '1' (meningioma). While the model




Fig. 3. Training and Validation Accuracy for MobileNet Baseline Model (60 Epochs).




Fig. 4. Training and Validation Loss for Xception Baseline Model (50 Epochs).

predicts the correct class, the confidence score is relatively low (49.03%). This uncertainty in classical predictions, even when correct, motivates exploring alternative architectures like the proposed hybrid QCNN, which might provide more robust feature representations or decision boundaries.

### C. Analysis Summary

The classical transfer learning models establish a strong performance benchmark on the processed data subset (~96-98% validation accuracy). Xception showed particularly good training stability and achieved the highest accuracy on this subset. However, the validation set size (52 images) limits the ability to draw definitive conclusions about generalization or to clearly distinguish between these powerful architectures. The results validate the data preprocessing pipeline and the classical feature extraction approach, while instances of low




Fig. 5. Example Prediction by Classical VGG16 Model on Image from 1.mat.

prediction confidence suggest potential room for improvement via quantum-enhanced methods.

### VIII. DISCUSSION

The results presented demonstrate the successful establishment of high-performing classical baselines for brain tumor classification on the prepared MRI dataset subset. The transfer learning approach using VGG16, MobileNetV1, and Xception achieved validation accuracies exceeding 96%, confirming the effectiveness of these architectures as feature extractors for this task. These results serve as a critical reference point for evaluating any potential quantum enhancement.

The motivation for exploring the hybrid QCNN architecture stems from the hypothesis that quantum circuits, operating on classically extracted features, might capture complex patterns or correlations that are challenging for purely classical dense layers [5], [8]. The defined 4-qubit VQC (Section IV-B), with its angle embedding and entangling layers, provides a mechanism to test this hypothesis. By processing the 4-dimensional feature vector from the classical head, the quantum layer could potentially transform the feature space in a non-classical way, possibly leading to improved separation of tumor classes, especially for ambiguous cases like the one shown in Fig. 5.

Metaphorically, while classical CNNs act as powerful "experts" on the high-dimensional image data, the quantum circuit could become a specialized "expert" on the compressed, abstract feature representation provided by the classical network. This differs from architectures like D-MoE where experts explicitly target different input conditions (e.g., SNR); here, the specialization is implicitly defined by the stage of processing within the hybrid pipeline.

However, significant challenges remain, primarily stemming from the limitations discussed further below. The small dataset subset used for these initial results necessitates validation on the full dataset. Furthermore, the true advantage of a small (4-qubit) quantum circuit integrated with a deep classical network is an

active area of research [10]. Issues like the choice of encoding, ansatz expressivity versus trainability (barren plateaus [11]), and the overhead of quantum simulation or noisy execution on real hardware must be carefully considered during the implementation and evaluation of the full hybrid system. The deprecation of the specific 'KerasLayer' also requires adapting to current PennyLane integration techniques.

### IX. CHALLENGES AND LIMITATIONS

This study, in its current preparatory stage, faces several limitations. The primary limitation is the use of a small dataset subset (198 training, 52 validation images derived from 250 .mat files). This significantly restricts the statistical power of the validation results and may lead to overly optimistic performance estimates that might not generalize to the full dataset (e.g., 3064 images mentioned in literature [12]). Accessing, processing, and evaluating on the complete dataset is crucial for drawing robust conclusions about both the classical baselines and the potential benefits of the quantum-enhanced model. The high, plateauing validation accuracy also signals the need for more rigorous evaluation protocols, such as k-fold cross-validation or testing on independent, external datasets, to better assess generalization.

Regarding the quantum component, the proposed 4 qubit circuit is relatively small, constrained by current simulation capabilities and near-term hardware limitations. While suitable for initial exploration on NISQ devices or simulators, its capacity to provide a significant advantage for a complex classification task already well-handled by deep classical feature extractors remains an open research question [5], [10]. Key challenges for the future hybrid implementation include selecting the optimal classical-to-quantum data encoding strategy (beyond simple angle embedding), designing an effective variational ansatz that avoids issues like barren plateaus [11], and managing the computational overhead of quantum circuit simulation or execution. Furthermore, the practicalities of running on real, noisy quantum hardware (if attempted) introduce additional complexities related to noise mitigation and error correction. Finally, the noted deprecation of the `qml.qnn.KerasLayer` necessitates adapting the implementation to use current best practices for integrating PennyLane QNodes with TensorFlow/Keras.

### X. CONCLUSION AND FUTURE WORK

This paper presented the essential groundwork for investigating a quantum-enhanced approach to brain tumor classification using MRI data. A systematic data preprocessing pipeline was developed and implemented, transforming raw .mat files into a format suitable for deep learning frameworks. Strong classical performance baselines were established using transfer learning with VGG-16, MobileNetV1, and Xception architectures on a preliminary dataset subset. These models achieved high validation accuracies (~96-98%), serving as crucial benchmarks, although the limitations imposed by the small validation set size were noted (Section IX). Concurrently, the core components of a 4 qubit hybrid QCNN were defined



using PennyLane, including the variational quantum circuit and its integration interface (`qml.qnn.KerasLayer`), outlining a concrete plan for quantum integration.

The immediate future work involves implementing and training the complete hybrid QCNN architecture, as conceptualized in Section IV. This requires integrating the defined quantum layer (using updated PennyLane methods if necessary) with the pre-trained classical feature extractors and performing end-to-end optimization. A critical next step is to apply this process to the full, larger brain tumor dataset to enable robust performance evaluation and comparison against the classical baselines. Evaluation will employ a comprehensive set of metrics (e.g., accuracy, precision, recall, F1-score, AUC) and potentially cross-validation to rigorously assess any advantages offered by the hybrid approach. Further research directions include exploring alternative quantum circuit designs (ansatzes and encoding strategies), investigating the impact of varying qubit numbers and circuit depth, analyzing the model's performance on challenging or ambiguous samples, and potentially assessing robustness to simulated noise to gauge feasibility for NISQ hardware. The ultimate goal is to determine empirically whether integrating quantum processing provides tangible benefits for this critical medical image analysis task.

#### ACKNOWLEDGMENT

The authors wish to express their sincere gratitude to their project supervisor, Dr. P. Dhana Lakshmi, for her invaluable guidance and support throughout this research project. We also acknowledge the creators of the brain tumor dataset used in this study [12] and thank Google Colab for providing computational resources.

#### REFERENCES

- [1] I. Cong, S. Choi, and M. D. Lukin, "Quantum convolutional neural networks," *Nature Physics*, vol. 15, no. 12, pp. 1273-1278, 2019.
- [2] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195-202, 2017.
- [3] M. Schuld, I. Sinayskiy, and F. Petruccione, "An introduction to quantum machine learning," *Contemporary Physics*, vol. 56, no. 2, pp. 172-185, 2015.
- [4] M. Henderson, S. Shakyia, S. Pradhan, and T. Cook, "Quantum convolutional neural networks: Powering image recognition with quantum circuits," *Quantum Machine Intelligence*, vol. 2, no. 1, pp. 1-9, 2020.
- [5] A. Mari, T. R. Bromley, J. Izaac, M. Schuld, and N. Killoran, "Transfer learning in hybrid classical-quantum neural networks," *Quantum*, vol. 4, p. 340, 2020.
- [6] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," *Physical Review A*, vol. 98, no. 3, p. 032309, 2018.
- [7] E. Grant, M. Benedetti, S. Cao, A. Hallam, J. Lockhart, V. Stojevic, et al., "Hierarchical quantum classifiers," *npj Quantum Information*, vol. 4, no. 1, p. 65, 2018.
- [8] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Physical Review Letters*, vol. 122, no. 4, p. 040504, 2019.
- [9] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, et al., "Noisy intermediate-scale quantum (NISQ) algorithms," *Reviews of Modern Physics*, vol. 94, no. 1, p. 015004, 2022.
- [10] C. Wilson, J. Szyman'ski, and M. F. Gonzalez-Zalba, "Optimizing the depth of variational quantum algorithms is NP-hard," *Quantum*, vol. 5, p. 463, 2021.
- [11] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, "Barren plateaus in quantum neural network training landscapes," *Nature Communications*, vol. 9, no. 1, p. 4812, 2018.
- [12] J. Cheng, W. Huang, S. Cao, R. Yang, W. Yang, et al., "Enhanced performance of brain tumor classification via tumor region augmentation and partition," *PloS one*, vol. 10, no. 10, p. e0140381, 2015. [Online]. Available: <https://doi.org/10.1371/journal.pone.0140381> [Dataset available at: <https://doi.org/10.6084/m9.figshare.1512427.v5>]
- [13] Z. N. K. Swati, Q. Zhao, M. Kabir, F. Ali, Z. Ali, et al., "Brain tumor classification for MR images using transfer learning and fine-tuning," *Computerized Medical Imaging and Graphics*, vol. 75, pp. 34-46, 2019.
- [14] S. Deepak and P. M. Ameer, "Brain tumor classification using deep CNN features via transfer learning," *Computers in Biology and Medicine*, vol. 111, p. 103345, 2019.
- [15] A. M. Ismael and A. S. Engu'r, "Deep learning approaches for brain tumor types classification based on MRI images," *Expert Systems with Applications*, vol. 165, p. 114042, 2021.
- [16] N. Ghassemi, A. Shoeibi, and M. Rouhani, "Deep neural network for brain tumor detection using magnetic resonance images," in *2020 International Conference on Machine Vision and Image Processing (MVIP)*, 2020, pp. 1-5.
- [17] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, M. S. Alam, et al., "PennyLane: Automatic differentiation of hybrid quantum-classical computations," *arXiv preprint arXiv:1811.04968*, 2018. [Online]. Available: <https://pennylane.ai>
- [18] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, et al., "TensorFlow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016. [Online]. Available: <https://www.tensorflow.org>
- [19] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [20] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "MobileNets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016, pp. 770-778.
- [22] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2017, pp. 1251-1258.
- [23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [24] E. Farhi and H. Neven, "Classification with quantum neural networks on near term processors," *arXiv preprint arXiv:1802.06002*, 2018.
- [25] J. P. Meyer, O. Kyriienko, and G. Quiroz, "Quantum computing in medical imaging: A review," *Medical Physics*, vol. 49, no. 11, pp. 7286-7306, 2022.
- [26] P. Tiwari, L. Qian, Q. Li, B. Wang, A. K. Singh, et al., "Quantum machine learning in medical image analysis: A review," *Computers in Biology and Medicine*, vol. 163, p. 107114, 2023.