Quantum Transfer Learning with Pretrained CNN Features for Multi-Class Medical Image Classification

Furkan Atban[1,2][0000-0002-1712-5155], Muhammed Yusuf Küçükkara[1,2][0000-0003-0600-3651], and Cüneyt Bayılmış[3][0000-0003-1058-7100]

[1] Department of Computer Engineering, Faculty of Technology, Sakarya University of Applied Sciences, Sakarya, Turkiye furkanatban@subu.edu.tr, muhammedkucukkara@subu.edu.tr

[2] Department of Computer Engineering, Institute of Natural Sciences, Sakarya University, Sakarya, Turkiye

[3] Department of Computer Engineering, Faculty of Computer and Information Sciences, Sakarya University, Sakarya, Turkiye cbayilmis@sakarya.edu.tr

Abstract. The automated classification of Optical Coherence Tomography (OCT) images is critical for the early diagnosis of retinal diseases such as choroidal neovascularization, diabetic macular edema and drusen. Classical deep networks achieve high accuracy but often require millions of parameters and significant computational resources. We present a hybrid Quantum Transfer Learning (QTL) framework combining a fine-tuned ResNet34 encoder with a variational quantum circuit (VQC) classifier for multi-class OCT classification. The classical backbone is partially fine-tuned—its first two layers remain frozen while the third and fourth blocks are adapted to OCT textures—and a classical projection reduces the 512-dimensional feature vector to a 5-dimensional input for a VQC acting on five qubits. Hyperparameters such as qubit count and circuit depth are selected via Particle Swarm Optimization (PSO). Experiments on the OCT2017 dataset show that the proposed model achieves 89.38% accuracy, supporting the feasibility of a parameter-efficient hybrid alternative for medical image classification under NISQ-inspired constraints. The results support the feasibility of hybrid quantum–classical models for medical imaging in the Noisy Intermediate-Scale Quantum era.

Keywords: Quantum Transfer Learning · Medical Image Classification · Optical Coherence Tomography · Hybrid Neural Networks · Particle Swarm Optimization

1 Introduction

Optical Coherence Tomography (OCT) produces high-resolution cross-sectional retinal images and is widely used in clinics for early diagnosis and monitoring of retinal diseases. Manual interpretation of OCT volumes is time-consuming and subject to inter-observer variability, motivating automatic classification systems. Convolutional neural networks (CNNs) and transfer learning achieve state-of-the-art results on medical images[1–3], yet their large parameter counts and training cost may limit deployment. Recent advances in quantum computing and quantum machine learning suggest that hybrid quantum–classical models could offer parameter-efficient alternatives[4–6].

As emphasised in ophthalmology literature, the early detection of retinal pathologies such as choroidal neovascularisation, diabetic macular oedema and drusen is vital for preventing vision loss. OCT provides micrometre-resolution cross-sections of retinal layers, allowing clinicians to visualise subtle structural changes. However, analysing large OCT volumes manually places a heavy workload on experts and may introduce subjective bias. Automated interpretation based on deep learning seeks to assist clinicians by rapidly screening for disease patterns with high sensitivity and specificity, while maintaining reproducibility across hospitals.

Traditional CNNs excel at feature extraction but typically involve tens of millions of learnable parameters and require significant GPU resources for training and inference. Such computational demands may limit their adoption in clinical settings where hardware is scarce or rapid deployment is needed. Moreover, purely classical models are sometimes criticised for limited flexibility in modelling decision boundaries under class imbalance or small dataset regimes. These constraints motivate the exploration of hybrid quantum–classical strategies that aim to reduce the number of trainable parameters while preserving classification accuracy.

In parallel, the field of quantum computing has witnessed rapid progress, yet current devices operate in the Noisy Intermediate-Scale Quantum (NISQ) regime with limited qubits and coherence times. Hybrid approaches that offload heavy feature extraction to classical networks and use shallow quantum circuits for decision making represent a feasible pathway towards near-term quantum advantage. In this context, Quantum Transfer Learning (QTL) leverages pretrained classical encoders to generate compact feature vectors that are subsequently processed by variational quantum circuits (VQCs). The integration of classical and quantum modules yields a learning pipeline that is both hardware-aware and adaptable to real-world medical imaging tasks.

Hybrid schemes leverage classical networks for feature extraction while delegating the decision stage to a variational quantum circuit (VQC). In Quantum Transfer Learning (QTL), a pretrained CNN provides high-level feature vectors that are encoded into a quantum state and processed by a parametrised quantum circuit[7, 8]. The circuit parameters are updated via classical optimisation, reflecting the variational quantum algorithm approach[9]. Current NISQ devices have limited qubit counts and are susceptible to noise; therefore, hybrid architectures offer a feasible intermediate solution between model expressivity and hardware limitations.

This paper proposes a QTL approach tailored for multi-class OCT classification. We partially fine-tune a ResNet34 backbone on the OCT2017 dataset, project its 512-dimensional feature vectors onto a five-dimensional subspace, and classify them with a depth-12 VQC acting on five qubits. Particle Swarm Optimization (PSO) selects optimal values for qubit count, circuit depth and other hyperparameters. The contributions are: (i) demonstrating a hybrid CNN–VQC architecture with partial fine-tuning for OCT classification, (ii) introducing PSO-based hyperparameter selection for quantum circuits, and (iii) reporting experimental results that evaluate the feasibility of hybrid classification under constrained quantum resources.

The remainder of the paper is organised as follows. Section 2 summarizes related hybrid quantum–classical approaches. Section 3 describes the pretrained CNN, VQC design, PSO optimisation and the full QTL pipeline. Section 4 details the dataset, evaluation metrics and experimental setup. Section 5 presents the results and comparisons. Conclusions and future research directions are discussed in Section 6.

2 Related Work

Hybrid quantum–classical learning combines classical representation learning with quantum decision making to address the limitations of NISQ devices. Early studies explored variational quantum classifiers operating on small classical feature sets[5]. Matic et al. evaluated several hybrid CNN–quantum classifiers for radiology and observed performance comparable to classical networks under controlled settings[10]. Azevedo et al. applied transfer learning with quantum layers to mammography, reporting improved generalisation[11]. Khan et al. proposed a hybrid quantum CNN for brain tumour classification, demonstrating high accuracy with fewer parameters[12]. Radhi et al. reviewed quantum machine learning in medical imaging and emphasised the need for standardised datasets and benchmarking protocols[13].

Subsequent works refined the quantum layer design. Ajlouni et al. introduced adaptive hybrid quantum CNNs with tunable entanglement structures for various imaging tasks[14]. Senokosov et al. compared parallel quantum circuits and quanvolution layers on MNIST and Medical MNIST, highlighting the potential of quantum feature mapping[15]. Jain et al. applied QTL to fundus images for diabetic retinopathy detection, providing preliminary evidence of quantum advantages on retinal data. While these studies illustrate promising directions, systematic evaluations on multi-class OCT classification remain limited. Our work addresses this gap by integrating a partially fine-tuned ResNet34 encoder with a VQC and systematically tuning quantum hyperparameters via PSO.

Beyond these pioneering efforts, several authors have examined how classical transfer learning can be seamlessly combined with variational quantum circuits. For instance, the work by Radhi et al. synthesises results from a broad range of applications—ranging from dermatological lesion analysis to volumetric MRI classification—and concludes that hybrid architectures often match classical accuracy with far fewer trainable weights. Ajlouni et al.'s adaptive quantum CNN explores dynamic entanglement patterns to adapt the quantum layer's expressive power to the complexity of the input data. Senokosov et al. demonstrate that even simple quanvolutional layers, which apply fixed quantum filters to image patches, can enrich feature representations when appended to classical CNNs. These studies collectively highlight that quantum circuits are versatile components that can be tailored to specific imaging modalities and that their integration into classical frameworks requires careful hyperparameter tuning. Despite these advances, the application of QTL to OCT data—where three disease classes and normal retina must be discriminated—remains largely unexplored. Therefore, our study contributes by systematically evaluating a hybrid CNN–VQC model on the OCT2017 benchmark with a rigorous PSO-based hyperparameter search.

3 Methodology

Our hybrid framework comprises a classical encoder, a variational quantum classifier and a PSO-driven hyperparameter search. Key components are summarised below.

3.1 Pretrained CNN and Feature Extraction

We employ a ResNet34 pretrained on ImageNet as the classical encoder. To adapt high-level representations to OCT data while maintaining general features, we freeze the initial two residual blocks and fine-tune the third and fourth blocks on the OCT2017 training set. For an input image $\mathbf{x}$, the encoder outputs a 512-dimensional feature vector $\mathbf{h} = \phi(\mathbf{x}; \psi)$ defined by Eq. (1). The network's class probabilities are obtained using the softmax function in Eq. (2).

$$\mathbf{h} = \phi(\mathbf{x}; \psi), \quad \mathbf{h} \in \mathbb{R}^{512} \tag{1}$$

$$\hat{\mathbf{p}} = \mathrm{softmax}(\mathbf{o}), \quad \hat{p}_c = \frac{\exp(o_c)}{\sum_{j=1}^{C} \exp(o_j)} \tag{2}$$

ResNet34 comprises an initial convolutional stem followed by four sequential residual blocks that facilitate the training of deep networks by introducing identity skip connections[2]. These residual connections allow gradients to flow directly to earlier layers, mitigating the vanishing gradient problem. In transfer learning scenarios, the early convolutional layers tend to extract generic edge and texture primitives, whereas deeper layers learn task-specific abstractions. By freezing the first two blocks, we preserve features learned from the large-scale ImageNet dataset, which are also relevant for OCT textures. Fine-tuning the third and fourth blocks allows the model to specialise its representations to the subtle morphological differences between CNV, DME, drusen and normal retina while keeping the parameter count modest.

3.2 Classical–Quantum Projection and VQC

The high-dimensional features are projected to a low-dimensional vector $\mathbf{z} = g(\mathbf{h};\omega)$ (Eq. (3)) using a learnable linear layer, reducing the 512 dimensions to n elements matching the number of qubits. Each component of $\mathbf{z}$ determines the rotation angle of a single-qubit gate. The variational quantum circuit, specified in Eq. (4), consists of angle embedding followed by layers of parameterised single-qubit rotations and entangling CNOT gates. The measurement of n Pauli-Z expectation values yields a real-valued vector $\mathbf{s}(\mathbf{z}) \in \mathbb{R}^n$. These expectations are mapped via a classical post-processing network to the final output.

$$\mathbf{z} = g(\mathbf{h};\omega), \quad \mathbf{z} \in \mathbb{R}^n \tag{3}$$

$$s_k(\mathbf{z}) = \left\langle 0^{\otimes n} \middle| U^\dagger(\mathbf{z},\theta)\, Z_k\, U(\mathbf{z},\theta) \middle| 0^{\otimes n} \right\rangle, \quad k = 1,\ldots,n \tag{4}$$

We adopt an angle-embedding strategy wherein each element of $\mathbf{z}$ is encoded into the rotation angle of a R_y gate acting on an individual qubit[5]. This encoding preserves the ordering of features and enables continuous parameterisation of qubit states. Following the encoding layer, we stack d layers of a strongly entangling ansatz composed of trainable R_y rotations and entangling CNOT gates arranged in a circular pattern. Entanglement is essential for modelling correlations that cannot be captured by separable qubit states, and the circuit depth d controls the expressive power of the quantum model. In our experiments, PSO selects $n = 5$ and $d = 12$ as optimal settings. After measurement, the resulting expectation values are passed to a classical dense layer with h hidden units (here $h = 166$) and a ReLU activation, followed by a softmax output. This post-processing stage allows the hybrid model to learn complex decision boundaries in classical space while keeping the quantum circuit shallow.

3.3 Particle Swarm Optimisation

Selecting the qubit count n, circuit depth d and learning rate η crucially impacts performance. We adopt Particle Swarm Optimisation (PSO) for data-driven hyperparameter tuning. Each particle encodes a candidate configuration $[n, d, \eta, q_\delta, h]$, where q_δ is a circuit initialisation scale and h denotes the number of hidden units in the post-processing network. Particle positions and velocities are updated according to Eqs. (5)–(6). The objective function, defined in Eq. 7, maximises validation accuracy over a small number of epochs. Table 1 lists the optimal hyperparameters obtained by PSO and used in final training.

$$\mathbf{v}_i^{t+1} = \omega\, \mathbf{v}_i^t + c_1 r_1(\mathbf{p}_i^t - \mathbf{x}_i^t) + c_2 r_2(\mathbf{g}^t - \mathbf{x}_i^t), \tag{5}$$
$$\mathbf{x}_i^{t+1} = \mathbf{x}_i^t + \mathbf{v}_i^{t+1} \tag{6}$$

$$\boldsymbol{\lambda}^* = \arg\max_{\boldsymbol{\lambda}\in\Omega}\ \mathrm{Acc}_{\mathrm{val}}(\boldsymbol{\lambda}) \tag{7}$$

Table 1. Best hyperparameters obtained via PSO (configuration used in the final training).

Hyperparameter	**Value**
Learning rate	0.002587
Quantum circuit depth (q_{depth})	12
Quantum step parameter (q_{δ})	0.0286
Number of hidden units	166
Batch size	60
Number of qubits (n_{qubits})	5

PSO is a population-based heuristic that simulates a flock of particles exploring the search space[16]. At each iteration, a particle moves towards a weighted combination of its own best-known position and the swarm's global best. The inertia term encourages exploration, while cognitive and social coefficients balance individual and collective learning. Unlike gradient-based methods, PSO does not require derivative information and can handle discrete or non-smooth objective landscapes. In our implementation, we initialise the swarm with random hyperparameter vectors drawn from a predefined range, evaluate the corresponding models on the validation set using a few training epochs, and iteratively update particle positions. After convergence, the best configuration is adopted for full training on the combined training and validation sets.

3.4 Quantum Transfer Learning (QTL)

Quantum Transfer Learning (QTL) extends the classical transfer learning approach to the domain of quantum machine learning by combining a pretrained classical model with a parametrised quantum circuit. In classical transfer learning, features learned on a source task are reused on a target task; QTL retains this idea but performs the final decision stage on a quantum device. The aim is to exploit the strong representational power of deep neural networks while capitalising on the expressive capacity of quantum circuits for low-dimensional decision boundaries. Heavy data processing and feature extraction remain on the classical side, whereas the quantum circuit handles the classification of a compact feature vector. This division of labour is particularly suitable for current NISQ devices, which offer a limited number of qubits and are sensitive to noise[17, 6].

In the QTL framework, the pretrained classical encoder is typically used as a fixed or partially fine-tuned feature extractor, preserving the knowledge acquired from large datasets and avoiding the need to train millions of parameters in tandem with the quantum circuit. The quantum component, usually a variational quantum circuit, serves as the learnable classifier. Its parameters are updated using classical optimisation routines, leading to a training loop closely resembling that of classical neural networks[5, 4]. This hybrid training scheme enables the seamless integration of quantum models into established deep learning workflows.

Concept and Definition The core idea behind QTL is to leverage a pre-trained classical network ϕ to map an input image $\mathbf{x}$ to a high-dimensional feature vector $\mathbf{h} = \phi(\mathbf{x}; \psi)$ and then use a quantum circuit as a classifier acting on a low-dimensional representation of $\mathbf{h}$. The classical network provides rich, transferable representations learned from large image datasets, while the quantum circuit offers a potentially more expressive decision function with a small number of trainable parameters. Because the classical features are reused, the quantum circuit can remain shallow—an important consideration for NISQ hardware. During training, only the parameters of the projection layer, the quantum circuit and the post-processing network are updated, keeping the early convolutional layers of the encoder fixed. This approach reduces the trainable parameter burden of the decision stage and may help mitigate overfitting when labelled medical data are scarce.

Classical-to-Quantum Feature Mapping A critical step in QTL is mapping the high-dimensional classical features $\mathbf{h} \in \mathbb{R}^{512}$ to a vector $\mathbf{z} \in \mathbb{R}^{n}$ compatible with an n-qubit quantum circuit. In our framework, this is accomplished by a learnable linear projection g as defined in Eq. (3). The projection reduces the dimensionality of the feature space from 512 to n (with $n = 5$ in our experiments) and is trained jointly with the quantum circuit. Although such a drastic reduction may induce information loss, the projection layer and the variational quantum circuit are trained jointly so that the compressed representation can still retain discriminative information relevant to the final classification task. After projection, each element of $\mathbf{z}$ is encoded into the quantum circuit via angle embedding, where the value determines the rotation angle of a single-qubit R_y gate. This embedding procedure allows continuous classical values to influence quantum states directly. The choice of projection dimension n and the specific encoding strategy are design decisions that influence the balance between hardware constraints and model capacity[5].

The classical–quantum feature mapping thus functions as a bridge between the extensive feature representations learned by the CNN and the restricted qubit space of the quantum circuit. By optimising the projection layer and the quantum circuit parameters together, the QTL model adapts the classical features to the quantum domain while supporting effective decision formation through subsequent classical post-processing.

3.5 Quantum Transfer Learning Pipeline

Figure 1 illustrates the complete pipeline. An OCT image is passed through the partially fine-tuned ResNet34 to extract 512 features. These are linearly projected to a 5-dimensional vector (for $n = 5$ qubits) and encoded into the quantum circuit via angle embedding. The depth-12 VQC produces expectation values which are fed into a classical network with 166 hidden units and a softmax output. The entire projection layer, quantum circuit parameters and post-processing weights are trained end-to-end using the cross-entropy loss on the training set. Hyperparameters are fixed from the PSO stage to ensure reproducibility.

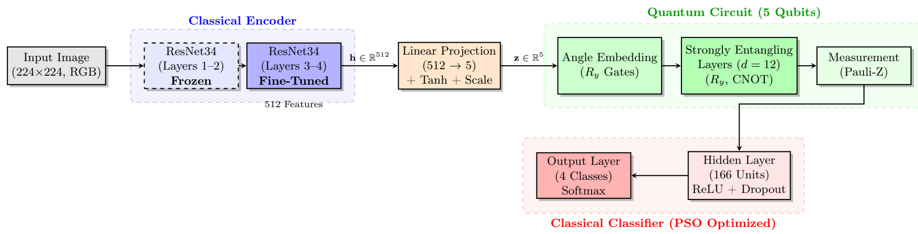

Fig. 1. The proposed hybrid architecture. The ResNet34 backbone is partially fine-tuned (Layers 3–4) to extract 512-dimensional features. These are projected to a 5-qubit quantum state, processed by a variational quantum circuit (depth=12), and classified by a classical dense network (166 hidden units) optimised via PSO.

4 Experiments

4.1 Dataset and Preprocessing

We evaluate the proposed QTL framework on the OCT2017 dataset, which contains OCT slices of four classes: CNV, DME, drusen and normal retina[3]. Images are resized to 224×224 pixels and converted to three-channel inputs by duplicating the grayscale image across RGB channels. To enhance generalisation, random rotations and horizontal flips are applied during training. We build a balanced subset with 200 images per class for training and validation, reserving a separate test set for performance assessment.

The OCT2017 dataset comprises over 109,000 B-scan slices drawn from 4,686 patients, with class distributions heavily skewed towards CNV and normal categories[3]. To manage the computational cost of quantum simulation while ensuring class balance, we sample 800 images (200 per class) for training and validation and another 800 images for testing. Prior to augmentation, each grayscale image is histogram normalised and padded to preserve aspect ratio before resizing to the required input dimensions. During training, we apply random 10-degree rotations and horizontal flips with probability 0.5 to mitigate overfitting. No augmentation is applied to the validation or test sets.

4.2 Experimental Setup

Training is performed on the TRUBA high-performance computing infrastructure using CUDA-accelerated GPUs for the classical network and the simulation of the quantum circuit. The quantum circuit is simulated under a noise-free assumption. Training and validation splits are fixed, and PSO uses only the validation subset. The performance metrics—accuracy, precision, recall and F1-score—are computed on the held-out test set using Eqs. (8)–(9). A confusion matrix visualises class-wise predictions.

$$\text{Accuracy} = \frac{\sum_{i=1}^{C} \text{TP}_i}{N} \tag{8}$$

$$\text{F1}_{\text{macro}} = \frac{1}{C} \sum_{i=1}^{C} \frac{2 \cdot \text{TP}_i}{2 \cdot \text{TP}_i + \text{FP}_i + \text{FN}_i} \tag{9}$$

Our experiments utilise the PyTorch framework for classical computations and PennyLane for quantum circuit simulation. Each model is trained using the Adam optimiser with the learning rate η selected by PSO and a mini-batch size of 60. Early stopping on the validation accuracy prevents overfitting. The classical encoder's fine-tuned layers and the projection, quantum circuit and post-processing parameters are trained jointly for 30 epochs using cross-entropy loss. Quantum simulations are performed in an idealised noise-free environment; investigating the effect of realistic noise models and error-mitigation techniques is left for future work.

5 Results

This section presents the experimental results of the proposed Quantum Transfer Learning (QTL) approach on the OCT2017 dataset. All results were obtained on the test dataset that was kept completely independent from the training and hyperparameter optimization processes. The performance metrics of the final experiments performed using the optimal configuration determined by PSO-based hyperparameter optimization are reported below.

The proposed QTL model achieved 89.38% accuracy on the test set. This result shows that the hybrid classical–quantum approach provides a feasible and stable performance under the evaluated experimental setting for the multi-class medical image classification problem. The macro-averaged precision, recall and F1-score were calculated as 0.8948, 0.8938 and 0.8929, respectively. These metrics indicate that the model exhibits a balanced performance across classes and does not overfit to a single class.

When the class-wise performance analysis is examined, it is seen that the highest F1-score (0.9417) is obtained for the NORMAL class. This indicates that the model can distinguish healthy retina images with high accuracy. The F1-scores for the CNV and DME classes were calculated as 0.8822 and 0.8889, respectively, indicating high and consistent classification performance for these classes as well. In the DRUSEN class, however, the recall is relatively lower (0.8050), while the precision is high (0.9200). This result shows that the model makes more cautious decisions in the DRUSEN class and limits false positive predictions.

The class-wise precision, recall and F1-score values obtained are presented in Table 2. The results show that the model generally exhibits balanced performance across classes; however, the DRUSEN class has relatively low recall.

Table 2. Class-wise performance results on the test set.

Class	Precision	Recall	F1-score	Support
CNV	0.8844	0.8800	0.8822	200
DME	0.8598	0.9200	0.8889	200
DRUSEN	0.9200	0.8050	0.8587	200
NORMAL	0.9151	0.9700	0.9417	200
Accuracy			0.8938	800
Macro avg	0.8948	0.8938	0.8929	800
Weighted avg	0.8948	0.8938	0.8929	800

Table 3. Comparison of the proposed hybrid QTL approach with state-of-the-art studies.

Study	Method	Accuracy (%)
Kermany et al. [3]	InceptionV3 (Transfer Learning)	96.60
Lee et al. [18]	VGG16 (Deep Learning)	93.45
Rong et al. [19]	Surrogate-assisted CNN	94.20
Proposed Method	**Hybrid ResNet34 + VQC (Ours)**	**89.38**

The class-wise decision behaviour of the model was examined in more detail via the confusion matrix presented in Fig. 2. A closer inspection shows that the NORMAL class is the most separable category, with the lowest error rate among all classes. The DME and CNV classes also exhibit strong classification performance, although a limited degree of confusion is observed between them, which may be attributed to partially overlapping retinal patterns. In contrast, DRUSEN remains the most challenging class, consistent with its lower recall reported in Table 2. Overall, the confusion structure indicates that most errors are concentrated among clinically similar disease categories, whereas healthy retina images remain comparatively easier to distinguish.

The PSO-based hyperparameter optimization determined that 5 qubits, a quantum circuit depth of 12, and a learning rate of approximately 0.0026 provide a balanced structure between the expressive power of the variational quantum circuit and training stability. The results obtained under this configuration show that the quantum classifier component functions effectively as a decision maker on the classical features extracted by the pretrained CNN.

As seen in Table 3, classical deep learning models still achieve higher classification accuracy on OCT2017. Nevertheless, the proposed hybrid model operates under a highly compressed 512-to-5 feature interface and a limited quantum parameter space. Therefore, its contribution should be interpreted as a parameter-efficient proof-of-concept for low-dimensional hybrid decision modeling under constrained quantum resources, rather than as evidence of superiority over high-capacity classical CNNs. The results indicate that a fine-tuned ResNet34 backbone coupled with a variational quantum circuit can preserve discriminative

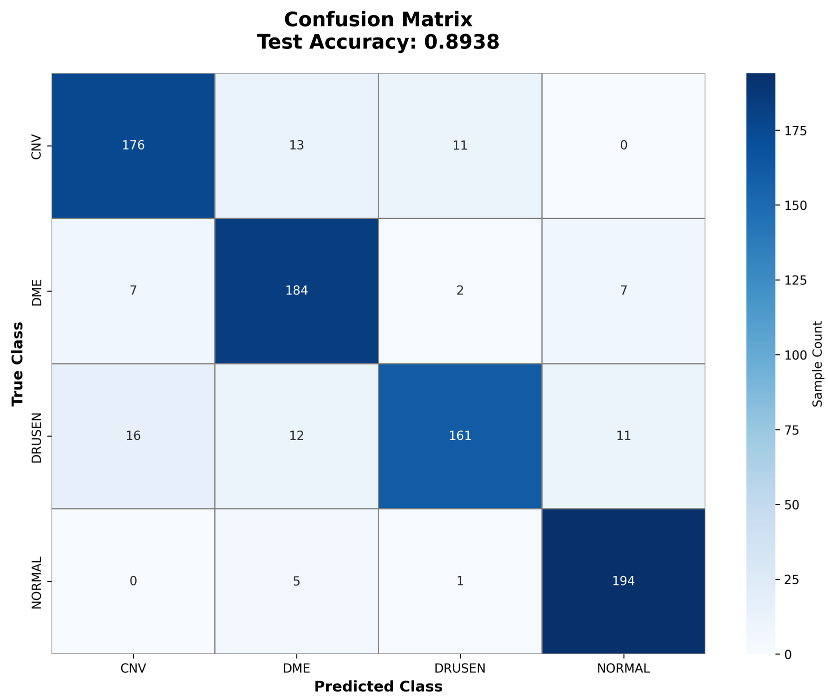

Fig. 2. Confusion matrix of the proposed QTL model on the independent test set. Each class contains 200 samples, enabling direct inspection of class-wise error distributions and absolute misclassification counts.

class information under severe dimensional and quantum resource constraints. In this respect, the study supports the feasibility of quantum transfer learning for medical imaging in the NISQ-inspired setting considered here.

Overall, the experimental results obtained show that the proposed Quantum Transfer Learning approach provides a feasible, balanced and interpretable performance for the multi-class medical image classification problem. When the quantitative performance metrics and confusion matrix analysis are evaluated together, it can be seen that the hybrid classical–quantum architecture makes a meaningful contribution to the classification process.

6 Conclusion

In this study, a hybrid Quantum Transfer Learning (QTL) approach that brings together classical and quantum learning components for the multi-class medical image classification problem has been proposed. The proposed method aims to leverage the powerful representation extraction capacity of a pretrained convolutional neural network and combine it with a variational quantum circuit, in order to exploit quantum learning models at the decision stage.

Experimental results, based on evaluations performed on the OCT2017 dataset, show that the proposed QTL approach provides a feasible and stable proof-of-concept for the multi-class medical image classification problem. The results obtained under the quantum circuit configuration determined by PSO-based hyperparameter optimization indicate that the quantum classifier component can operate effectively on the deep features extracted by the classical CNN within the evaluated experimental setting. Class-wise performance metrics and confusion matrix analysis further show that the model can draw balanced distinctions among classes, while most errors remain concentrated among clinically similar categories.

An important advantage of the proposed approach is the reduction of the parameter count by freezing the first two layers of the ResNet34 backbone and fine-tuning the deeper layers (Layers 3–4), while allowing the quantum circuit to be trained in a smaller, controllable parameter space. This structure provides a useful framework for studying hybrid quantum–classical models under near-term quantum constraints, particularly in settings where parameter efficiency and low-dimensional decision modeling are of interest. Moreover, PSO-based hyperparameter selection allowed for a more systematic model configuration by reducing experimental bias in quantum circuit design.

However, in this study, quantum circuit simulations were conducted under the assumption of an ideal and noise-free environment. The effects of noise and errors that may arise on real quantum hardware were left outside the scope of this work. Future studies plan to evaluate the proposed QTL approach on noisy quantum circuits, to investigate different quantum ansatz structures, and to conduct comprehensive comparisons with larger datasets. In addition, the applicability of the quantum component to different clinical imaging modalities is regarded as an important direction to be investigated in the future.

In conclusion, this study demonstrates that hybrid approaches based on the combined use of classical deep learning and quantum learning models offer a promising research direction for medical image classification problems. The findings support that the Quantum Transfer Learning approach provides a viable and open framework for research in the context of near-term quantum technologies.

Acknowledgments. This research was supported by the TUBITAK 2211/C National PhD Scholarship Program in the Priority Fields in Science and Technology entitled "Improving the Performance of Machine Learning Algorithms with Quantum Computing".

The numerical calculations reported in this paper were partially performed at TUBITAK ULAKBIM, High Performance and Grid Computing Center (TRUBA resources).

References

1. Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.

2. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
3. Daniel S Kermany, Michael Goldbaum, Wenjia Cai, Carolina CS Valentim, Huiying Liang, Sally L Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, et al. Identifying medical diagnoses and treatable diseases by image-based deep learning. *cell*, 172(5):1122–1131, 2018.
4. Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.
5. Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Physical review letters*, 122(4):040504, 2019.
6. John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.
7. Andrea Mari, Thomas R Bromley, Josh Izaac, Maria Schuld, and Nathan Killoran. Transfer learning in hybrid classical-quantum neural networks. *Quantum*, 4:340, 2020.
8. Amira Abbas, David Sutter, Christa Zoufal, Aurélien Lucchi, Alessio Figalli, and Stefan Woerner. The power of quantum neural networks. *Nature Computational Science*, 1(6):403–409, 2021.
9. Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016.
10. A. Matic, Maureen Monnet, J. Lorenz, Balthasar Schachtner, and Thomas Messerer. Quantum-classical convolutional neural networks in radiological image classification. pages 56–66, 09 2022.
11. Vanda Azevedo, Carla Silva, and Inês Dutra. Quantum transfer learning for breast cancer detection. *Quantum Machine Intelligence*, 4(1):5, 2022.
12. Muhammad Al-Zafar Khan, Abdullah Al Omar Galib, Nouhaila Innan, and Mohamed Bennai. Brain tumor diagnosis using quantum convolutional neural networks. *arXiv preprint arXiv:2401.15804*, 2024.
13. Eman A Radhi, Mohammed Y Kamil, and Mazin Abed Mohammed. Quantum machine and deep learning for medical image classification: A systematic review of trends, methodologies, and future directions. *Iraqi Journal for Computer Science and Mathematics*, 6(2):9, 2025.
14. Naim Ajlouni, Adem Özyavaş, Mustafa Takaoğlu, Faruk Takaoğlu, and Firas Ajlouni. Medical image diagnosis based on adaptive hybrid quantum cnn. *Research Square (Research Square)*, 06 2023.
15. Arsenii Senokosov, Alexandr Sedykh, Asel Sagingalieva, Basil Kyriacou, and Alexey Melnikov. Quantum machine learning for image classification. *Machine Learning Science and Technology*, 5:15040–15040, 02 2024.
16. James Kennedy and Russell Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95-international conference on neural networks*, volume 4, pages 1942–1948. ieee, 1995.
17. Maria Schuld and Francesco Petruccione. Supervised learning with quantum computers. *Quantum science and technology*, 17, 2018.
18. Cecilia S Lee, Doug M Baughman, and Aaron Y Lee. Deep learning is effective for classifying normal versus age-related macular degeneration oct images. *Ophthalmology Retina*, 1(4):322–327, 2017.

19. Yibiao Rong, Dehui Xiang, Weifang Zhu, Kai Yu, Fei Shi, Zhun Fan, and Xinjian Chen. Surrogate-assisted retinal oct image classification based on convolutional neural networks. *IEEE journal of biomedical and health informatics*, 23(1):253–263, 2018.