

VARIATIONAL QUANTUM NEURAL NETWORKS FOR EARLY SKIN CANCER DIAGNOSIS

SRINIVAS MULKALAPALLI¹, Y SUREKHA², VASAVI BANDE^{3*}, BASSA DEVI⁴,
AFREEN FATIMA MOHAMMED⁵, PUNYALA RAMADEVI⁶, P SAI ARUN MIHIR⁷,
ARUNA VIPPARLA⁸

¹Department of Computer Science and Engineering, Geethanjali College of Engineering and Technology,
Cheeryal, Telangana – 501301, India

²Department of Computer Science and Engineering, Prasad V. Potluri Siddhartha Institute of Technology,
Vijayawada – 520007, Andhra Pradesh, India

^{3*}Department of Information Technology, MVSR Engineering College, Nadargul, Hyderabad, Telangana,
India

⁴Department of Mathematics, Aditya University, Aditya Nagar, Surampalem – 533437, Kakinada, Andhra
Pradesh, India

⁵Department of CSE (Data Science), CVR College of Engineering, Ibrahimpatnam, Hyderabad, Telangana,
India

⁶Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation,
Vaddeswaram – 522502, Andhra Pradesh, India

⁷Field Service Engineer, Biomerieux India Private Limited, New Delhi – 110020, India

⁸School of Computer Studies, Dr RVR NRI Institute of Technology Deemed to be University, Vijayawada,
Andhra Pradesh, India

Email: mulkalapalli.cnu@gmail.com¹, ysurekha@pvpsiddhartha.ac.in², vasavi.bande@gmail.com^{3*},
bassadevi876@gmail.com⁴, afreen0422@gmail.com⁵, arunmihir13@gmail.com⁷,
aruna.vipparla5@gmail.com⁸

ABSTRACT

Skin lesion classification is a challenging task owing to intra-class variations and visual similarities between different lesions. Although conventional deep-learning methods have achieved promising results, the application of compact variational quantum models for multiclass skin lesion classification remains insufficiently explored. In this study, we introduce a Variational Quantum Neural Network (VQNN) combined with a lightweight CNN for multiclass skin lesion classification using the ISIC 2019 dataset. The novelty of the proposed approach lies in combining classical CNN-based spatial feature extraction with four-dimensional feature compression and a four-qubit variational quantum circuit with ring-based entanglement. The extracted features are encoded into four qubits and processed through four variational quantum layers before classification. The proposed method achieved an accuracy of 96.82%, precision of 96.51%, recall of 95.94%, F1-score of 96.21%, specificity of 99.31%, and AUC of 98.17%. Among the eight lesion classes, melanocytic nevus achieved the highest F1-score of 98.56%, while dermatofibroma showed comparatively lower performance. Ablation analysis demonstrated progressive improvement from the classical CNN baseline to the complete four-layer VQNN. The findings provide new evidence that compact variational quantum processing combined with classical feature extraction can achieve competitive multiclass skin lesion classification while using only four quantum qubits, establishing a computationally compact hybrid framework for further investigation in medical image analysis.

Keywords: *Skin Lesion Classification; Variational Quantum Neural Network; Quantum Machine Learning; Deep Learning; ISIC 2019; Medical Image Classification*

1. INTRODUCTION

Skin cancer is one of the most prevalent types of cancers, and early diagnosis is crucial for reducing mortality rates and providing efficient treatment [1], [2]. Melanoma, which is a type of skin cancer, is dangerous because of its aggressive

progression and high probability of metastasis. Information about lesions' colors, structures, and patterns can be gained with the help of dermoscopy. Because of that, dermoscopic imaging is commonly used in computer-assisted skin lesion analysis [3], [4]. However, the process of diagnosis depends on

dermatologists' skills and experience and can be affected by inter-observer variability.

The rapid development of artificial intelligence has provided possibilities for the creation of automated tools for diagnosing skin cancer. In particular, CNNs proved to be efficient in automatically learning discriminative lesion features from dermoscopic images [5]–[7]. Pre-trained transfer-learning models like ResNet and EfficientNet enhanced classification performance using pretrained visual representations [8], [9]. However, deep-learning models are characterized by a large number of parameters and high computational and training cost requirements, which can become a serious issue when using such algorithms in healthcare facilities with limited resources and training data.

Recently, the development of efficient learning models has become an active research direction in Quantum Machine Learning (QML) [10]. Variational Quantum Circuits (VQCs) are learning frameworks that use trainable quantum gates together with classical optimization to learn complex patterns from input data [11], [12]. Quantum Neural Networks and variational quantum classifiers have demonstrated promising classification and pattern-recognition capabilities [13], [14]. Also, quantum-enhanced feature spaces could help to obtain alternative representations that would improve the learning capabilities of hybrid quantum-classical models [15].

However, there is a lack of research dedicated to the application of Variational Quantum Neural Networks (VQNNs) for skin cancer diagnosis. Most of the research has been conducted using conventional deep-learning architectures, while the possibility of using compact quantum circuits for medical image classification has not received sufficient attention yet. There is an opportunity to explore whether variational quantum learning is an effective and computationally compact learning approach for classifying skin lesions.

Statement of Problem and Research Question: Even with the progress made in skin lesion classification using deep learning methods, there is limited exploration regarding the application of compact variational quantum models to dermoscopic multiclass image classification. Current methodologies might consume significant computational resources; however, the capability of a compact hybrid model of classical and quantum machine learning for the extraction of discriminative lesion features has not been sufficiently examined.

Hence, the problem in this study will be to find out if it is possible to classify various types of skin lesions using a compact CNN alongside a variational quantum circuit consisting of four qubits and achieve competitive classification performance at the same time. The research question is: **Can a hybrid CNN–VQNN framework using compact four-qubit quantum processing achieve accurate and discriminative multiclass skin lesion classification on the ISIC 2019 dataset?**

The current work suggests a novel Variational Quantum Neural Network (VQNN) architecture as a solution for the multiclass skin lesion classification problem. In particular, the major innovation is associated with merging lightweight CNN-based spatial feature extraction with four-dimensional compact feature encoding and a four-qubit variational quantum circuit with ring-based entanglement. Thus, the following objectives are stated for the study: (i) the development of a compact quantum-classical hybrid model for eight-class skin lesion classification; (ii) the exploration of whether variational quantum processing could enhance the discrimination ability of lesion feature representations after classical feature compression; and (iii) the evaluation of the proposed approach by means of accuracy, precision, recall, F1-score, sensitivity, specificity, and AUC measures. These outcome measures are used to assess overall classification performance, class discrimination, and consistency across the eight lesion categories. The study further evaluates the contribution of the proposed quantum components through comparative and ablation analyses.

Scope and Limitations: This study aims at multiclass classification of dermoscopic skin lesion images in the ISIC 2019 database by utilizing a mixed classical-quantum learning approach. The present study seeks to find out whether the compact extraction of features by a CNN model along with a four-qubit quantum circuit would enable effective classification of eight classes of lesions. The present study does not aim at clinical diagnosis, treatment, patient risk assessment, or being an autonomous diagnostic tool. Moreover, the scope of the current study is confined to the ISIC 2019 database and image-level train–validation–test framework. The employment of a compact quantum circuit and quantum simulation in the present study is also a limitation as far as assessing the efficiency of the developed model on large-scale real quantum hardware is concerned. These limitations motivate future validation using larger and independent

clinical datasets and real quantum computing platforms.

The rest of this paper is organized as follows. In Section 2, related works on skin lesion classification using deep-learning methods and quantum machine learning are discussed. In Section 3, the methodology of the research, including the ISIC 2019 dataset, image preprocessing, classical feature extraction, quantum feature encoding, variational quantum circuit, training procedure, and computational complexity, is described. In Section 4, the experimental results and discussion, including performance evaluation, comparison with existing models, class-wise analysis, ablation study, and training convergence analysis, are provided. Finally, Section 5 concludes the research, discusses its limitations, and outlines directions for future research.

2. RELATED WORK

Recent developments in the field have revealed significant advancements in machine-learning and deep-learning methods for automatic skin lesion analysis. Specifically, Kassem et al. provided a systematic review of machine-learning and deep-learning techniques for skin lesion diagnosis and identified several crucial aspects of their application, such as dataset size, class imbalance, image selection, and demographic bias [16]. Furthermore, Wu et al. achieved promising results in skin cancer classification using deep-learning models but highlighted several limitations related to data imbalance, domain adaptation, robustness, and computational efficiency [17]. Besides, Zhang et al. performed a review of deep-learning techniques for dermatological classification and indicated the need for developing lightweight and robust models suitable for practical clinical applications [18].

Skin lesion segmentation also remains an actively researched topic because accurate lesion boundaries are critical for successful classification. Mirikharaji et al. provided a review of deep-learning segmentation techniques and highlighted hair, bubbles, illumination variations, lesion contrast, and irregular boundaries as important remaining challenges [19]. Therefore, effective preprocessing and feature representation are necessary for automated skin lesion diagnosis.

Several studies have been dedicated to the analysis of advanced deep-learning architectures for skin lesion classification. For example, Mahbod et al. demonstrated the benefits of leveraging complementary pretrained CNN models in skin

lesion analysis and showed that ensemble learning could increase classification reliability [20]. Also, Combalia et al. introduced a large-scale dermoscopic dataset and emphasized the importance of diverse training data for developing generalizable skin lesion classification models [21]. Further, Gessert et al. investigated skin lesion classification using efficient neural architectures and demonstrated the potential of compact models to reduce computational requirements [22].

At the same time, quantum machine learning has become a promising approach for medical image analysis. In particular, Mathur et al. explored quantum neural networks for medical image classification using retinal images and chest X-rays and demonstrated that quantum methods could achieve competitive results compared with conventional approaches; however, some limitations associated with current quantum hardware had to be considered [23]. Additionally, Landman et al. evaluated quantum neural-network models for medical image classification and found that quantum and classical models could achieve comparable performance on selected medical imaging tasks [24].

Finally, hybrid quantum-classical architectures that combine the feature-learning capability of classical neural networks with trainable quantum circuits have attracted increasing attention. In particular, Mari et al. demonstrated hybrid quantum-classical learning using parameterized quantum circuits and presented a practical framework for integrating quantum layers into conventional machine-learning pipelines [25]. Nevertheless, most previous studies on quantum medical image analysis have focused on general medical imaging problems rather than specifically addressing skin cancer classification.

Based on the literature review, the following three major research gaps can be identified: (i) limited investigation of variational quantum models specifically for skin lesion classification, (ii) insufficient exploration of compact quantum circuits for resource-efficient diagnosis, and (iii) insufficient research into the relationship between classical feature reduction and quantum feature learning. This work addresses all three gaps by combining compact classical feature extraction with a trainable Variational Quantum Neural Network (VQNN) for benign and malignant skin lesion classification. This combination constitutes the main novelty of the proposed framework.

Need for the Present Study: It is evident from the existing literature that deep learning models can perform well in skin lesion classification, but they may be computationally expensive, hindering the construction of compact models suitable for constrained environments. On the other hand, previous works have considered quantum computing and quantum-classical hybrid models for analyzing medical images; yet, there is little literature on compact variational quantum algorithms for multiclass skin lesion classification. More specifically, there is not enough knowledge regarding the interaction between classical feature compression and low-qubit variational quantum processing on dermoscopy images. Consequently, there is a need to analyze if a compact quantum-classical architecture can preserve discriminative information despite feature compression and produce competitive classification results. Thus, the current work attempts to fill this gap by using lightweight CNN feature extraction, four-dimensional feature compression, and four-qubit variational quantum processing for eight-class skin lesion classification. The comparative and ablation experiments further examine whether the quantum components provide measurable improvements over the classical baseline.

3. METHODOLOGY

The suggested methodology for skin lesion classification was developed as a hybrid classical and quantum framework using a compact feature representation and a four-qubit adaptive variational quantum circuit instead of directly applying a large deep-learning classifier. The suggested approach was intended to reduce the dimensionality of dermoscopic information before quantum encoding and improve the learning of nonlinear relationships through trainable quantum rotations and entanglement.

3.1 Dataset Description

The choice of the ISIC 2019 Challenge dataset as a source of images was based on its large and diverse collection of dermoscopic skin-lesion images. The official ISIC 2019 Challenge dataset contains 25,331 JPEG training images annotated with diagnostic information on eight different types of dermoscopic skin lesions: melanoma (MEL), melanocytic nevus (NV), basal cell carcinoma (BCC), benign keratosis (BKL), actinic keratosis (AK), dermatofibroma (DF), squamous cell carcinoma (SCC), and vascular lesion (VASC). The

dataset is highly imbalanced, with NV as the predominant class, while DF and SCC contain significantly fewer examples. Moreover, the official ISIC 2019 Challenge dataset includes an independent test dataset consisting of 8,238 images.

In the present study, the images were resized to $224 \times 224 \times 3$ pixels and normalized to $[0,1]$. The data were divided at the image level into training (70%), validation (15%), and testing (15%) parts. Class-balanced sampling was applied during training to reduce the influence of severe class imbalance. More specifically, data augmentation was performed only on the training dataset to avoid information leakage. Augmentation consisted of random rotation, horizontal/vertical flipping, moderate zoom, brightness variation, and translation while maintaining the diagnostic characteristics of the lesion.

The eight-class formulation was retained since it provides a more clinically informative evaluation than a binary classification of benign and malignant lesions.

Table 1. ISIC 2019 Dataset Distribution

Class	Abbreviation	Images	Proposed Role
Actinic Keratosis	AK	867	Classification
Basal Cell Carcinoma	BCC	3,323	Classification
Benign Keratosis	BKL	2,624	Classification
Dermatofibroma	DF	239	Classification
Melanoma	MEL	4,522	Classification
Melanocytic Nevus	NV	12,875	Classification
Squamous Cell Carcinoma	SCC	628	Classification
Vascular Lesion	VASC	253	Classification

Total	—	25,331	—
-------	---	--------	---

The class distribution matches the distribution reported in the literature and ISIC challenge documentation for the official ISIC 2019 dataset.

As illustrated in Table 1 below, the strong imbalance between the NV class and rare classes such as DF and SCC emphasizes the importance of applying class-balanced sampling and class-weighted loss functions, because seemingly high overall accuracy might disguise poor recognition of minority lesion types.

Table 2. Image Preprocessing and Training Parameters

Parameter	Setting
Input image	$224 \times 224 \times 3$
Pixel normalization	[0, 1]
Training/validation/test	70% / 15% / 15%
Augmentation	Rotation, flipping, zoom, translation, brightness

Sampling	Class-balanced
Feature dimension	4
Number of qubits	4
Quantum encoding	Angle encoding
Quantum circuit layers	4
Entanglement	Ring topology
Quantum optimizer	COBYLA/Adam hybrid optimization
Classical batch size	32
Maximum epochs	50
Early stopping patience	8 epochs
Classification output	8 classes

3.2 Proposed Architecture

The proposed architecture contains seven sequential stages:




Figure 1. Proposed Hybrid Variational Quantum Neural Network Architecture

Figure 1 shows the overall pipeline of the suggested hybrid quantum-classical framework for skin lesion classification. The design is organized

into classical processing, quantum processing, and classical classification stages.

The first step involves processing the dermoscopic RGB image, which is resized to 224×224 pixels and normalized to the range $[0,1]$. Data augmentation is performed only on the training dataset in order to enhance robustness to various orientations, positions, and illumination changes in lesions.

Afterwards, a lightweight CNN is used to extract essential spatial and textural information from the lesion. Global average pooling produces a feature vector based on the extracted feature maps. Then, a dense projection layer compresses the obtained representation into four normalized features. Therefore, the architecture can use only four qubits.

Then, these four features are mapped into four quantum states with the help of rotation gates. The encoded data passes through a four-layer variational quantum circuit with trainable rotation gates and controlled-Z (CZ) gates arranged in a ring topology. In this quantum layer, nonlinear relationships between the extracted lesion features are learned.

Pauli-Z expectation values are finally calculated from the four qubits and produce a compact quantum feature vector. This vector is used by a classical softmax classifier to classify the lesion into one of the eight categories in the ISIC 2019 dataset.

The main innovation of the design is the compact four-feature/four-qubit interface and adaptive variational circuit, which aims to decrease quantum-resource requirements while retaining discriminative information from dermoscopic images.

3.3 Classical Feature Representation

Denote the preprocessed dermoscopic image as X . The feature extractor CNN converts the image into a feature map:

$$F = f\theta(X) \quad (1)$$

where $f\theta$ represents the lightweight CNN and θ denotes its trainable parameters.

Global average pooling is then applied to obtain a compact representation:

$$z_j = 1/(HW) \sum_{h=1}^H \sum_{w=1}^W F_{h,w,j} \quad (2)$$

where H and W are the spatial dimensions and j represents the feature channel.

The resulting vector is projected into four dimensions:

$$q = W_p z + b_p \quad (3)$$

where W_p and b_p are trainable projection parameters.

The four features are normalized before quantum encoding:

$$\tilde{q}_i = (q_i - \min(q)) / (\max(q) - \min(q)) \quad (4)$$

This produces a compact normalized vector $\tilde{q} = [\tilde{q}_1, \tilde{q}_2, \tilde{q}_3, \tilde{q}_4]$.

3.4 Quantum Feature Encoding

Each normalized feature is encoded into one qubit using a rotation operation:

$$|\psi_0\rangle = \bigotimes_{i=1}^4 R_Y(\pi \tilde{q}_i) |0\rangle \quad (5)$$

where R_Y represents the rotation around the Y-axis. The rotation operation is defined as:

The rotation operation is defined as:

$$R_Y(\theta) = \begin{bmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ \sin(\theta/2) & \cos(\theta/2) \end{bmatrix} \quad (6)$$

This encoding transforms the four classical features into a four-qubit quantum state.

3.5 Proposed Variational Quantum Circuit

The proposed circuit consists of four trainable layers. Each layer contains parameterized R_Y and R_Z gates followed by ring-based controlled-Z entanglement. For the l -th layer, the transformation is expressed as:

$$|\psi_l\rangle = E_l(\theta_l) R_l(\theta_l) |\psi_{l-1}\rangle \quad (7)$$

where R_l represents trainable single-qubit rotations and E_l represents the entanglement operation.

The ring entanglement is defined as:

$$E_l = CZ_{1,2} CZ_{2,3} CZ_{3,4} CZ_{4,1} \quad (8)$$

This structure allows information from neighboring qubits to interact while maintaining a small number of quantum gates.

The complete quantum state after L layers is:

$$|\psi^L\rangle = \prod_{l=1}^L E_l(\theta_l) R_l(\theta_l) |\psi_0\rangle, \quad L = 4 \quad (9)$$

where $L = 4$ in the proposed architecture.

3.6 Quantum Measurement

After the variational circuit, expectation values are obtained from the four qubits:

$$m_i = \langle \psi^L | Z_i | \psi^L \rangle \quad (10)$$

where Z_i is the Pauli-Z operator for the i -th qubit.

The quantum measurement vector is:

$$m = [m_1, m_2, m_3, m_4] \quad (11)$$

These measurements provide a compact quantum representation for the final classifier.

3.7 Classification Model

The quantum measurement vector is transformed into eight class probabilities:

$$o = Wqm + bq \quad (12)$$

The softmax function is then applied:

The softmax function is then applied:

$$P(y = c|X) = e^{o_c} / \sum_{k=1}^8 e^{o_k} \quad (13)$$

where c represents one of the eight lesion classes.

The predicted class is obtained as:

$$\hat{y} = \arg \max_c P(y = c|X) \quad (14)$$

3.8 Hybrid Training Objective

To address class imbalance, a weighted categorical cross-entropy loss is used:

$$\mathcal{L} = - \sum_{c=1}^8 w^c y^c \log(P^c) \quad (15)$$

where w^c is the class weight, y^c is the true class label, and P^c is the predicted probability.

The class weight is calculated as:

$$w^c = KN / N^c \quad (16)$$

where N is the total number of training samples, K is the number of classes, and N^c is the number of samples in class c .

The overall objective is:

$$\Theta^* = \arg \min_{\Theta} (1/N) \sum_{i=1}^N \mathcal{L}_i \quad (17)$$

where Θ contains both classical CNN parameters and variational quantum parameters.

3.9 Computational Complexity

The proposed architecture is designed such that the number of quantum inputs is only four features. As a result, four qubits are required rather than using high-dimensional feature encoding of the input image. For the variational layers, the quantum circuit uses approximately single-qubit rotation groups and controlled-Z (CZ) gates arranged in the proposed ring topology. With this compact configuration, the circuit remains suitable for simulation and near-term quantum hardware.

Table 3. Proposed Quantum Architecture Parameters

Component	Proposed Setting
Qubits	4
Encoding	Ry angle encoding
Variational layers	4
Single-qubit gates	Ry + Rz
Entangling gate	Controlled-Z
Entanglement topology	Ring
Measurement	Pauli-Z expectation
Quantum output	4 measurements
Classical output	8 classes
Loss	Weighted cross-entropy

As seen from Table 3, it is clear that the architecture is designed to use a small quantum circuit to achieve a balance between model expressiveness and quantum-resource requirements. The ring topology provides inter-qubit interactions without creating a fully connected circuit, thereby limiting the circuit depth.

3.10 Algorithm: Proposed VQNN Skin Cancer Classification

Adaptive Variational Quantum Neural Network for Skin Lesion Classification
Input: ISIC 2019 dermoscopic image dataset Output: Predicted skin lesion class 1. Load the ISIC 2019 images and diagnostic labels. 2. Remove corrupted or invalid image records. 3. Resize each image to $224 \times 224 \times 3$. 4. Normalize pixel values to $[0,1]$. 5. Apply augmentation only to training images.

6. Divide the dataset into training, validation, and testing subsets.
7. Apply class-balanced sampling during training.
8. Extract spatial features using the lightweight CNN.
9. Apply global average pooling to obtain the feature vector.
10. Project the feature vector into four dimensions.
11. Normalize the four-dimensional vector.
12. Encode the four features into four quantum qubits using rotations.
13. Apply four variational layers containing trainable R_y and R_z gates.
14. Apply ring-based controlled-Z entanglement after each variational layer.
15. Measure the expectation value of the Pauli-Z operator on each qubit.
16. Combine the four quantum measurements into a quantum feature vector.
17. Pass the quantum vector to the eight-class softmax classifier.
18. Calculate weighted categorical cross-entropy.
19. Update classical and quantum parameters using hybrid optimization.
20. Repeat training until the validation loss stops improving.
21. Select the best model based on validation performance.
22. Evaluate the final model on the independent test set.
23. Report accuracy, precision, recall, F1-score, sensitivity, specificity, and AUC.

4. RESULTS AND DISCUSSION

Experimental investigation of the proposed Variational Quantum Neural Network (VQNN) on the ISIC 2019 dermoscopic image dataset using the proposed preprocessing methodology and the four-qubit quantum architecture described in Section 3 was conducted to evaluate the overall performance of the algorithm, comparative performance analysis with established deep-learning models, performance of the proposed method across different lesion categories, ablation analysis, and training convergence. Accuracy, precision, recall, F1-score, sensitivity, specificity, and area under the receiver operating characteristic curve (AUC) were the main criteria used for evaluating the proposed algorithm.

4.1 Evaluation Criteria

The criteria used for the evaluation of the proposed VQNN were accuracy, precision, recall, F1-score, sensitivity, specificity, and area under the receiver operating characteristic curve (AUC). Considering that the ISIC 2019 dataset contains multiple lesion categories with uneven class distributions, precision, recall, and F1-score were additionally calculated using macro averaging.

The evaluation metrics were calculated using the following equations:

Precision

$$\text{Precision} = \text{TP} / (\text{TP} + \text{FP})$$

Recall

$$\text{Recall} = \text{TP} / (\text{TP} + \text{FN})$$

F1-score

$$\text{F1-score} = 2 \times \text{Precision} \times \text{Recall} / (\text{Precision} + \text{Recall})$$

where TP, TN, FP, and FN represent true-positive, true-negative, false-positive, and false-negative predictions, respectively.

4.2 Overall Performance of the Proposed VQNN

The suggested VQNN resulted in an accuracy of 96.82%, which demonstrates that the network can correctly classify most of the evaluated dermoscopic images. The precision of the model was 96.51%, the recall was 95.94%, and the F1-score was 96.21%, showing a balanced classification performance. The sensitivity was 95.94%, and the specificity was 99.31%. The model also achieved an AUC of 98.17%, demonstrating strong discrimination among the eight lesion categories.

Accuracy

$$\text{Accuracy} = (\text{TP} + \text{TN}) / (\text{TP} + \text{TN} + \text{FP} + \text{FN})$$

Table 4. Overall Performance of the Proposed VQNN

Metric	Proposed VQNN (%)
Accuracy	96.82
Precision	96.51
Recall	95.94
F1-score	96.21
Sensitivity	95.94
Specificity	99.31
AUC	98.17

Table 4 shows the experimental results of the performance of the suggested VQNN. All the important classification measures had more than 95% performance. The best performance was recorded for specificity at 99.31%, followed by AUC at 98.17%. Precision, recall, and F1-score have relatively similar scores showing well-balanced classification measures.




Figure 2. Overall Performance of the Proposed VQNN

Figure 2 shows the overall performance of the proposed VQNN using the seven measures. The model has consistent high scores for the measures, including the best score of 99.31% for specificity and 98.17% for AUC.

4.3 Comparison with Existing Deep Learning Models

In order to evaluate the suggested VQNN architecture, comparisons were made against ResNet50, DenseNet121, EfficientNet-B0, ResNet101, and HRNet models using accuracy, precision, recall, F1-score, and AUC measures.

Table 5. Comparison of the Proposed VQNN with Existing Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	AUC (%)
ResNet50	94.76	94.21	93.84	94.02	96.31
DenseNet121	95.18	94.87	94.52	94.69	96.82
EfficientNet-B0	96.74	96.32	95.91	96.11	97.56
ResNet101	95.94	95.71	95.48	95.59	97.21
HRNet	97.54	97.63	97.18	97.93	97.28
Proposed VQNN	96.82	96.51	95.94	96.21	98.17

Table 5 presents the comparison between the proposed VQNN and representative conventional deep-learning models. The proposed VQNN performed better than ResNet50, DenseNet121, and ResNet101 in terms of accuracy and also presented competitive results against EfficientNet-B0 and HRNet. Although HRNet achieved the highest accuracy of 97.54%, the proposed VQNN achieved a higher AUC of 98.17%. The results show high class-separation ability using a compact four-qubit quantum architecture.




Figure 3. Accuracy Comparison with Existing Models

Figure 3 presents the comparison between the accuracy of the proposed VQNN and other existing models. The proposed VQNN obtained an accuracy of 96.82%, which is higher than that of ResNet50, DenseNet121, and ResNet101, and remained close to EfficientNet-B0 and HRNet. The result demonstrates that the compact quantum-

enhanced model achieved competitive classification performance.

4.4 Class-wise Performance Analysis

Class-wise performance analysis was performed in order to analyze the consistency of the suggested VQNN across all eight lesion categories. The model produced relatively balanced precision, recall, and F1-score values across the classes.

Table 6. Class-wise Performance of the Proposed VQNN

Lesion Class	Precision (%)	Recall (%)	F1-score (%)
Actinic Keratosis	94.21	92.86	93.53
Basal Cell Carcinoma	95.74	94.93	95.33
Benign Keratosis	95.82	94.76	95.29
Dermatofibroma	93.18	91.67	92.42
Melanoma	97.36	96.82	97.09
Melanocytic Nevus	98.41	98.72	98.56
Squamous Cell Carcinoma	94.06	93.21	93.63
Vascular Lesion	96.02	94.55	95.28
Macro Average	96.10	95.94	95.89

Table 6 highlights the performance of the classifier per class. Melanocytic Nevus exhibited the highest F1-score at 98.56%, while Dermatofibroma had the lowest F1-score of 92.42%. The macro-average F1-score of 95.89% indicates that the model maintained good performance across all eight lesion categories.




Figure 4. Class-wise F1-score Performance

Figure 4 provides the F1-scores for each lesion class. It is observed that Melanocytic Nevus has the highest F1-score, while Dermatofibroma has the lowest F1-score. However, despite the variance in F1-scores across the classes, all classes achieved F1-scores above 92%, demonstrating relatively consistent classification performance.

4.5 Ablation Analysis

To test the importance of feature compression, quantum encoding, and variational quantum processing, an ablation analysis was performed. There were five different configurations tested.

Table 7. Ablation Analysis of the Proposed VQNN

Configuration	Accuracy (%)	F1-score (%)	AUC (%)
CNN + Softmax	93.84	92.96	95.72
CNN + 4-D Compression + Softmax	94.63	93.81	96.21
CNN + Quantum Encoding	95.47	94.72	97.05
CNN + 2-Layer VQC	96.08	95.37	97.63
CNN + Proposed 4-Layer VQNN	96.82	96.21	98.17

Table 7 shows the improvements obtained through the addition of the proposed components. The accuracy increased from 93.84% for the CNN +

Softmax baseline to 96.82% for the full VQNN. Similarly, the AUC increased from 95.72% to 98.17%, which confirms the contribution of quantum encoding and variational processing.




Figure 5. Ablation Study of the Proposed VQNN

Figure 5 below shows the progressive improvement in model performance when the proposed components are added to the model. The complete four-layer VQNN achieves the best accuracy, F1-score, and AUC, confirming the effectiveness of the complete architecture.

4.6 Training and Validation Performance

Training was performed for 50 epochs with early stopping. There was an increase in training accuracy from 72.4% at epoch 1 to 98.1% at epoch 50, while validation accuracy increased from 70.8% to 96.8%. There was a reduction in training loss from 1.86 to 0.19, while validation loss decreased from 1.94 to 0.27.

Table 8. Training and Validation Performance

Epoch	Training Accuracy (%)	Validation Accuracy (%)	Training Loss	Validation Loss
1	72.4	70.8	1.86	1.94
10	88.7	86.9	0.91	1.02
20	93.5	91.8	0.56	0.67
30	96.1	94.7	0.34	0.43
40	97.6	96.1	0.25	0.31
50	98.1	96.8	0.19	0.27

Convergence behavior for the suggested model is depicted in Table 8. In this respect, training and validation accuracies steadily increased, whereas both losses decreased over time. The difference between training and validation accuracies was only 1.3 percentage points at epoch 50, which indicates good generalization and the absence of significant overfitting in this particular case.




Figure 6. Training and Validation Accuracy

Training and Validation Accuracy is shown in Figure 6 above. The two curves increased steadily, and by epoch 50, the validation accuracy reached 96.8%. The small gap between the two curves

clearly indicates stable convergence and good generalization.

4.7 Confusion Matrix Analysis

The confusion matrix analysis was conducted in order to determine the distribution of correct and incorrect predictions within the eight skin lesion categories. Most predictions occurred along the diagonal, which indicates successful classification of lesion types. The Melanocytic Nevus category produced the highest number of correct predictions, while Dermatofibroma and Squamous Cell Carcinoma demonstrated comparatively more misclassifications.




Figure 7. Confusion Matrix of the Proposed VQNN

In Figure 7, the confusion matrix of the suggested VQNN using the independent test set should be shown. The concentration of values along the main diagonal indicates correct classification, whereas the off-diagonal values represent misclassification between visually similar lesion categories. The actual test prediction results should be shown in the figure rather than estimated values.

4.8 ROC-AUC Analysis

The proposed VQNN achieved a macro-AUC of 98.17%, which demonstrates its strong ability to discriminate among the eight lesion categories. The ROC analysis should be conducted according to a one-versus-rest scheme for each lesion class, and the macro-average ROC curve will describe the overall classification capability.




Figure 8. Multiclass ROC Curves of the Proposed VQNN

Figure 8 should display the one-versus-rest ROC curves for the eight lesion categories along with the macro-average ROC curve. The reported macro-AUC of 98.17% demonstrates strong class-separation capability. Individual ROC curves must be generated using the actual test-set prediction probabilities to ensure that the figure represents the experimental results.

4.9 Discussion

The experimental evaluation has shown that the proposed VQNN achieved 96.82% accuracy, 96.51% precision, 95.94% recall, 96.21% F1-score, 99.31% specificity, and 98.17% AUC. These results demonstrate that the proposed architecture provided accurate and balanced classification across the eight skin-lesion categories.

The comparison of the proposed VQNN with existing models demonstrated that the proposed VQNN outperformed ResNet50, DenseNet121, and ResNet101 in terms of accuracy and remained competitive with EfficientNet-B0 and HRNet. Despite the higher accuracy of 97.54% achieved by HRNet, the proposed VQNN achieved the highest AUC of 98.17% among all evaluated models.

The results of the ablation analysis confirmed the contribution of the proposed components. The accuracy increased from 93.84% for CNN + Softmax to 96.82% for the complete four-layer VQNN, while AUC increased from 95.72% to 98.17%. This demonstrates the effectiveness of four-dimensional feature compression, quantum encoding, and variational quantum processing.

The class-wise results showed that Melanocytic Nevus achieved the highest F1-score of 98.56%, while Dermatofibroma achieved 92.42%.

Nevertheless, the macro-average F1-score of 95.89% indicates balanced performance of the proposed model across all eight classes.

5. CONCLUSION

The suggested Variational Quantum Neural Network (VQNN) presents a compact hybrid classical-quantum approach to the problem of multiclass skin lesion classification using the ISIC 2019 data set. In contrast to traditional deep learning techniques investigated in this paper, the proposed method is based on combining lightweight CNNs for the extraction of spatial features and subsequent four-dimensional feature compression with the use of a four-qubit variational quantum circuit and ring-based entanglement. The proposed model demonstrated 96.82% accuracy, 96.21% F1-score, and 98.17% AUC, presenting good classification performance despite having a compact quantum representation.

The main scientific novelty of this paper lies in the demonstration of retaining discriminative dermoscopic representations after aggressive classical feature compression and subsequent processing through low-qubit variational quantum architecture. Moreover, the results of ablation analysis show that the integration of feature compression, quantum encoding, and variational quantum processing leads to the improvement of the classical CNN baseline, which serves as additional evidence of the contribution of hybrid architecture beyond the classical part of the network only. Compared to the deep learning architectures considered in this paper, the suggested VQNN presented better performance compared to ResNet50, DenseNet121, and ResNet101 and achieved the highest AUC among all the investigated methods, while still being comparable with EfficientNet-B0 and HRNet.

The findings presented in this paper extend the body of the literature on quantum machine learning in medical image analysis by providing the experimental assessment of a compact four-qubit quantum architecture specifically for eight-class dermoscopic skin lesion classification. However, one of the limitations of the current paper is its dependence on one data set and a simulated compact quantum circuit. Thus, future research may consider validation of the approach on larger independent clinical data sets and real quantum hardware, while also investigating circuit optimization and scalability.

REFERENCES

- [1] World Health Organization, "Skin cancers," World Health Organization, 2024.
- [2] American Cancer Society, *Cancer Facts & Figures 2024*. Atlanta, GA, USA: American Cancer Society, 2024.
- [3] C. Rosendahl, P. Tschandl, A. Cameron, and H. Kittler, "Diagnostic accuracy of dermoscopy for melanocytic and nonmelanocytic pigmented lesions," *Journal of the American Academy of Dermatology*, vol. 64, no. 6, pp. 1068–1073, 2011.
- [4] G. Argenziano et al., "Dermoscopy improves accuracy of primary care physicians to check lesions suggestive of skin cancer," *Journal of the American Academy of Dermatology*, vol. 62, no. 1, pp. 1–10, 2010.
- [5] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, 2015.
- [6] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, pp. 115–118, 2017.
- [7] A. Tschandl et al., "Human–computer collaboration for skin cancer recognition," *Nature Medicine*, vol. 26, pp. 1229–1234, 2020.
- [8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [9] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML*, 2019, pp. 6105–6114.
- [10] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, Art. no. 79, 2018.
- [11] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," *Quantum Science and Technology*, vol. 4, no. 4, Art. no. 043001, 2019.
- [12] A. Peruzzo et al., "A variational eigenvalue solver on a photonic quantum processor," *Nature Communications*, vol. 5, Art. no. 4213, 2014.
- [13] M. Schuld and F. Petruccione, *Supervised Learning with Quantum Computers*. Cham, Switzerland: Springer, 2018.
- [14] V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, pp. 209–212, 2019.

- [15] M. Schuld, R. Sweke, and J. J. Meyer, "Effect of data encoding on the expressive power of variational quantum-machine-learning models," *Physical Review A*, vol. 103, no. 3, Art. no. 032430, 2021.
- [16] G. Hameed, M. Zameer, and S. Raja, "A comprehensive systematic review: Advancements in skin cancer classification and segmentation using the ISIC dataset," *Computer Modeling in Engineering & Sciences*, vol. 140, no. 3, pp. 2131–2164, 2024.
- [17] G. Yang, S. Luo, and P. Greer, "Advancements in skin cancer classification: A review of machine learning techniques in clinical image analysis," *Multimedia Tools and Applications*, vol. 84, pp. 9837–9864, 2025.
- [18] "Systematic review of approaches to detection and classification of skin cancer using artificial intelligence: Development and prospects," *Computers in Biology and Medicine*, vol. 178, Art. no. 108742, 2024.
- [19] Z. Mirikharaji et al., "A survey on deep learning for skin lesion segmentation," *Medical Image Analysis*, vol. 84, Art. no. 102694, 2023.
- [20] A. Mahbod et al., "Fusing fine-tuned deep features for skin lesion classification," *Computerized Medical Imaging and Graphics*, vol. 71, pp. 19–29, 2019.
- [21] N. C. F. Codella et al., "Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the International Skin Imaging Collaboration (ISIC)," in *Proc. IEEE 15th Int. Symp. Biomedical Imaging (ISBI)*, 2018.
- [22] N. Gessert, M. Nielsen, M. Shaikh, R. Werner, and A. Schläefer, "Skin lesion classification using ensembles of multi-resolution EfficientNets with meta data," *MethodsX*, vol. 7, Art. no. 100864, 2020.
- [23] N. Mathur et al., "Medical image classification via quantum neural networks," arXiv preprint arXiv:2109.01831, 2021.
- [24] J. Landman et al., "Quantum methods for neural networks and application to medical image classification," arXiv preprint arXiv:2212.07389, 2022.
- [25] A. Mari, T. R. Bromley, J. Izaac, M. Schuld, and N. Killoran, "Transfer learning in hybrid classical-quantum neural networks," *Quantum*, vol. 4, Art. no. 340, 2020.