WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

An Explainable Hybrid Quantum-Classical Deep Learning Framework for Early Skin Cancer Detection Using Compact Feature Representation

Bhargav Sai Vaddi [1], B. Sai Venkata Teja[2], Ch. Hari Prasad[3], D. Sai Teja[4]

(vaddibhargavsai@gmail.com)
(rudrateja112@gmail.com)
(hariprasadchandaka1824@gmail.Com)
(dattisaiteja072@gmail.com)

ABSTRACT :

Skin cancer is a major global health concern, and early detection is essential for improving treatment outcomes. This project proposes an explainable hybrid quantum-classical framework for skin cancer detection using the HAM10000 dermoscopic image dataset. ResNet-50 is employed as the deep learning backbone to extract meaningful features from skin lesion images. Three experimental configurations are investigated: direct classification using the ResNet-50 classification head, classification using extracted deep features, and classification using PCA-reduced features with 4, 6, and 8 dimensions. The extracted and reduced features are evaluated using classical machine learning algorithms, including Logistic Regression, Support Vector Machine, Random Forest, and XGBoost, alongside a PennyLane-based Variational Quantum Classifier (VQC). This comparison provides insights into the performance of classical and quantum classification approaches under different feature representations. To improve model interpretability, Grad-CAM generates visual heatmaps highlighting image regions contributing to predictions. The complete framework is integrated into a web application developed using React for the frontend and FastAPI for backend services. SQLite is used to maintain prediction and screening history. The application provides benign or malignant classification, malignancy probability scores, model information, and visual explanations, offering an integrated platform for comparative evaluation and explainable skin cancer screening.

Keywords: Skin Cancer Detection, ResNet-50, Principal Component Analysis (PCA), Variational Quantum Classifier (VQC), Grad-CAM

INTRODUCTION

Skin cancer is one of the most common cancers worldwide, and early detection is essential for improving treatment outcomes. Accurate skin lesion classification is challenging because different conditions may have similar visual characteristics. Although traditional diagnosis relies on dermatologists' clinical and dermoscopic examinations, Artificial Intelligence (AI), particularly deep learning and machine learning, offers promising solutions for automated skin lesion analysis.

This project proposes an Explainable Hybrid Quantum-Classical Framework for Skin Cancer Detection, combining deep feature extraction, dimensionality reduction, classical machine learning, and quantum machine learning. The framework uses the HAM10000 dataset, containing 10,015 dermoscopic images across seven diagnostic categories, mapped into benign and malignant classes. Images undergo resizing, normalization, and lesion-level splitting into training, validation, and testing sets to reduce data leakage between subsets.

ResNet-50 is used as the primary CNN architecture to extract high-level visual features, producing 2048-dimensional representations. Three experimental configurations are investigated: (1) direct classification using the ResNet-50 classification head, (2) classification using extracted deep features, and (3) classification using PCA-reduced features with 4, 6, and 8 dimensions. This enables comparison of different feature representations and classification strategies.

The extracted and reduced features are evaluated using four classical algorithms: Logistic Regression, Support Vector Machine (SVM), Random Forest, and XGBoost. Additionally, a Variational Quantum Classifier (VQC) is implemented using PennyLane. It uses parameterized quantum circuits with rotation and entanglement operations, and experiments are conducted with 4, 6, and 8 qubits using a quantum simulator.

To improve interpretability, Grad-CAM generates visual heatmaps highlighting image regions contributing to ResNet-50 predictions. The framework is integrated into a web application developed using React with Vite for the frontend and FastAPI for the backend. SQLite maintains screening and prediction history. The application displays classification results, malignancy probabilities, model information, inference details, and Grad-CAM visualizations.

Overall, the project covers dataset preparation, preprocessing, ResNet-50 feature extraction, dimensionality reduction, classical and quantum classification, comparative evaluation, explainability, and web application integration. It systematically investigates how different feature representations affect classification performance while exploring the integration of explainable AI and quantum machine learning into a unified skin cancer detection system.

LITERATURE SURVEY

The reviewed research papers provide the foundation for developing the proposed hybrid quantum-classical skin cancer detection framework. Previous studies demonstrate the effectiveness of CNN-based feature extraction, PCA-based dimensionality reduction, and quantum classifiers for skin lesion analysis. These studies also highlight the importance of comparing classical and quantum models and improving prediction interpretability using Grad-CAM. Inspired by these approaches, our project integrates ResNet-50, PCA, classical machine learning algorithms, and PennyLane-based VQC. These techniques are combined to evaluate different feature representations and classification configurations within a unified explainable framework.

A hybrid quantum-classical model was developed using QCNN and VQC with a compact 4-qubit circuit. PCA reduced image features to four components, while class-weighted loss addressed data imbalance. The model achieved 70.13% test accuracy, demonstrating the feasibility of compact quantum-based classification [1]A framework combining wavelet transformation, Quantum Chebyshev polynomials, and autoencoders was proposed for feature extraction and dimensionality reduction. LSTM was used for classification on ISIC-2017 and HAM10000 datasets. The model achieved 99.58% accuracy on HAM10000 [2].Quantum classifiers, including QSVC and QCNN, were compared with classical deep learning models such as ResNet50V2 and MobileNet. Using HAM10000, the QCNN achieved 82.86% accuracy, while MobileNet-assisted QSVC achieved 72.5%. The study explored the effectiveness of quantum classification methods [3].An explainable hybrid framework integrated DenseNet121 with an 8-qubit PennyLane classifier using parameterized quantum circuits. A soft-voting ensemble combined classical and quantum predictions, achieving 95.11% accuracy. SHAP, LIME, and Grad-CAM were incorporated to improve explainability [4].A hybrid framework combined image preprocessing, autoencoder-based hair removal, U-Net segmentation, and quantum classification. Extracted features were reduced and processed through a 4-qubit quantum circuit. The model achieved 99.67% classification accuracy on HAM10000 [5].A hybrid model combined MobileNetV2, BiLSTM, and quantum CNN for benign and malignant lesion classification. The quantum circuit used Hadamard and parameterized RY rotation gates. It achieved 89.3% test accuracy and 94.33% malignant-lesion recall [6].Q-SkinNet used ResNet-18 for feature extraction and Particle Swarm Optimization for feature selection. Six selected features were processed using a 6-qubit quantum circuit. The model achieved 98.2% accuracy on the ISIC-2019 dataset [7].DeepSkin was developed using DenseNet169 and ResNet50 with lesion segmentation and class-balancing techniques. Dull Razor filtering was applied to remove hair artifacts. DenseNet169 achieved 91.2% accuracy and a 91.7% F1-score on HAM10000 [8].A hybrid CNN framework combined DeepLabV3+ segmentation with features extracted using MobileNetV2, EfficientNetB0, and DenseNet201. ReliefF was used for feature selection, followed by KNN classification. The model achieved 94.42% accuracy on ISIC-2019 [9].A Swin Transformer-based approach was investigated using a dataset containing 41,975 dermoscopic images. Its shifted-window self-attention mechanism captures spatial relationships and image features. The model achieved 86.37% accuracy on an independent benchmark test set [10].A Swin Transformer-based framework combined lesion segmentation, multi-scale dilated convolutions, and attention mechanisms. It was evaluated on four skin lesion datasets and achieved 98.72% classification accuracy, demonstrating the effectiveness of transformer-based feature learning [11].An eight-layer CNN was developed using batch normalization, ReLU activation, and dropout. Random oversampling addressed class imbalance, while Grad-CAM provided visual explanations. The model achieved 98.50% accuracy on HAM10000 and 97.11% on ISIC-2019 [12].An ensemble framework combined ConvNeXt, SENet, DenseNet, and EfficientNet-B0 using weighted bootstrapping and data augmentation. Majority voting combined the model predictions. It achieved 96.12% cross-validation accuracy on HAM10000 [13].SkinNet-14, a compact transformer-based architecture, was introduced for low-resolution dermoscopic images. It uses convolutional tokenization and a transformer encoder to reduce computational requirements. It achieved 97.85% accuracy on HAM10000 with 2–8 seconds of training time per epoch [14].

ViTfSCD, a Vision Transformer, was developed using image patches, positional embeddings, and multi-head self-attention. The model was pretrained on ImageNet and fine-tuned on HAM10000. It achieved 94.1% accuracy and demonstrated performance on the external DERMOFIT dataset [15].

METHODOLOGY

1. Overview of the Approach

The proposed methodology begins with the HAM10000 dataset, containing 10,015 dermoscopic images, followed by image preprocessing and lesion-level data splitting. ResNet-50 is used to extract 2048-dimensional deep features from the processed images. Three experimental configurations are investigated: direct classification using the ResNet-50 classification head, classification using extracted features, and classification using PCA-reduced features of 4, 6, and 8 dimensions. Classical classifiers, including Logistic Regression, SVM, Random Forest, and XGBoost, are compared with a PennyLane-based Variational Quantum Classifier (VQC). Grad-CAM generates visual explanations of ResNet-50 predictions. Finally, the complete framework is integrated into a React and FastAPI web application with SQLite for maintaining prediction history. The system provides benign or malignant predictions, malignancy probabilities, model information, and visual explanations.

2.1. Proposed System Architecture

The proposed system consists of multiple stages, including image preprocessing, deep feature extraction, dimensionality reduction, classification, and explainability. ResNet-50 extracts high-level image features, which are evaluated using classical machine learning models and a Variational Quantum Classifier (VQC). Grad-CAM provides visual explanations, while a React and FastAPI-based application enables users to upload images and obtain predictions.

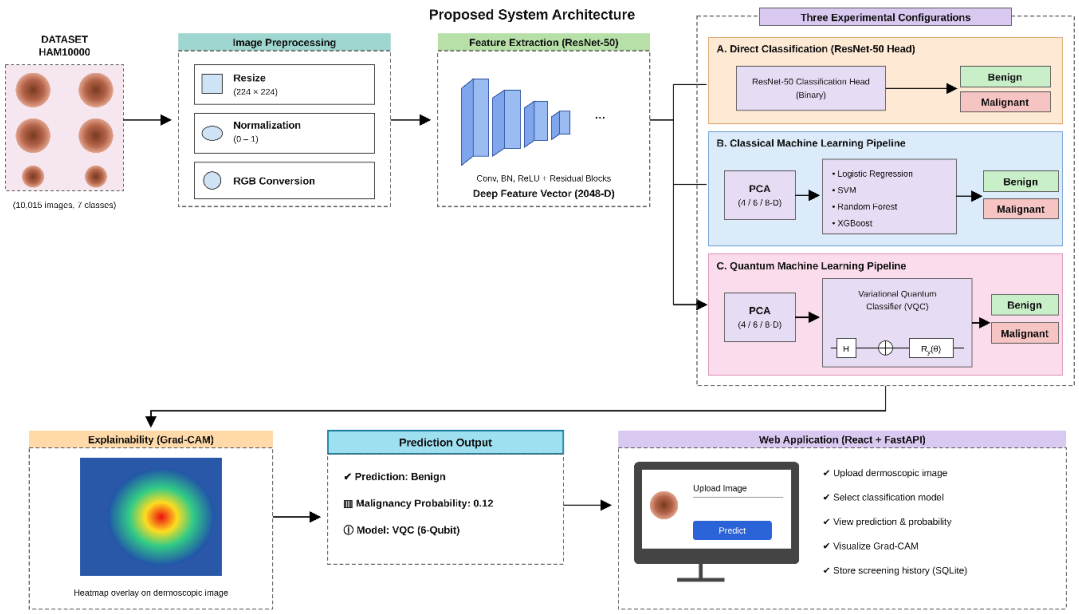


Figure 2.1: Overall Proposed System Architecture

2.2. Dataset Description

The HAM10000 dataset is used to develop and evaluate the proposed skin cancer detection framework. It contains 10,015 dermoscopic images belonging to seven diagnostic categories. These categories are grouped into benign and malignant classes to formulate a binary classification problem. The dataset is divided into training, validation, and testing subsets using lesion-level splitting to reduce data leakage.

2.3. Image Preprocessing

Image preprocessing prepares dermoscopic images for consistent feature extraction. The images are resized to 224 × 224 pixels and normalized according to the ResNet-50 input requirements. The processed images are then supplied to the feature extraction stage to obtain consistent visual representations.

2.4. Deep Feature Extraction Using ResNet-50

ResNet-50 is employed as the primary deep learning architecture for extracting meaningful visual features from skin lesion images. Its residual connections help train deeper networks by improving information flow across layers. The final feature representation contains 2048 dimensions and is used in subsequent classification experiments.

2.5. Feature Dimensionality Reduction Using PCA

Principal Component Analysis (PCA) is applied to reduce the dimensionality of the extracted deep features while retaining important information. The 2048-dimensional representations are transformed into compact feature spaces of 4, 6, and 8 dimensions. These reduced representations are used to investigate the effect of feature dimensionality on classical and quantum classification performance.

2.6. Classical Machine Learning Models

Four classical machine learning algorithms are used for skin lesion classification: Logistic Regression, Support Vector Machine (SVM), Random Forest, and XGBoost. These models are evaluated using extracted deep features and PCA-reduced representations. Their performance is compared to understand the effectiveness of different classification techniques.

2.7. Variational Quantum Classifier

A Variational Quantum Classifier (VQC) is implemented using the PennyLane framework. The input features are encoded into parameterized quantum circuits using rotation gates, while CNOT gates introduce entanglement between qubits. Experiments are conducted using 4, 6, and 8 qubits to evaluate different circuit configurations. The quantum models are executed using the PennyLane simulator.

2.8. Explainability Using Grad-CAM

Gradient-weighted Class Activation Mapping (Grad-CAM) is incorporated to improve the interpretability of the ResNet-50 predictions. It generates heatmaps highlighting image regions that contribute to the model's classification. These visual explanations help users understand which regions receive greater attention during prediction.

2.9. Web Application Implementation

The proposed framework is integrated into a web application using React with Vite for the frontend and FastAPI for backend services. Users can upload skin lesion images, select classification models, and view prediction results. SQLite maintains screening history, while the application displays malignancy probability scores, model information, and Grad-CAM visualizations.

3. IMPLEMENTATION AND RESULTS

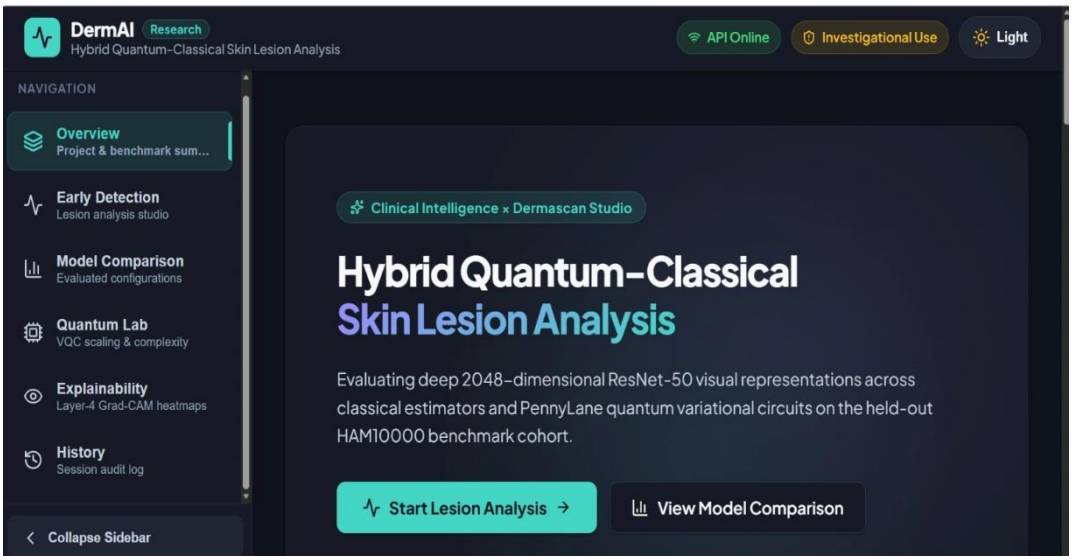


Figure 3.1: Home Page

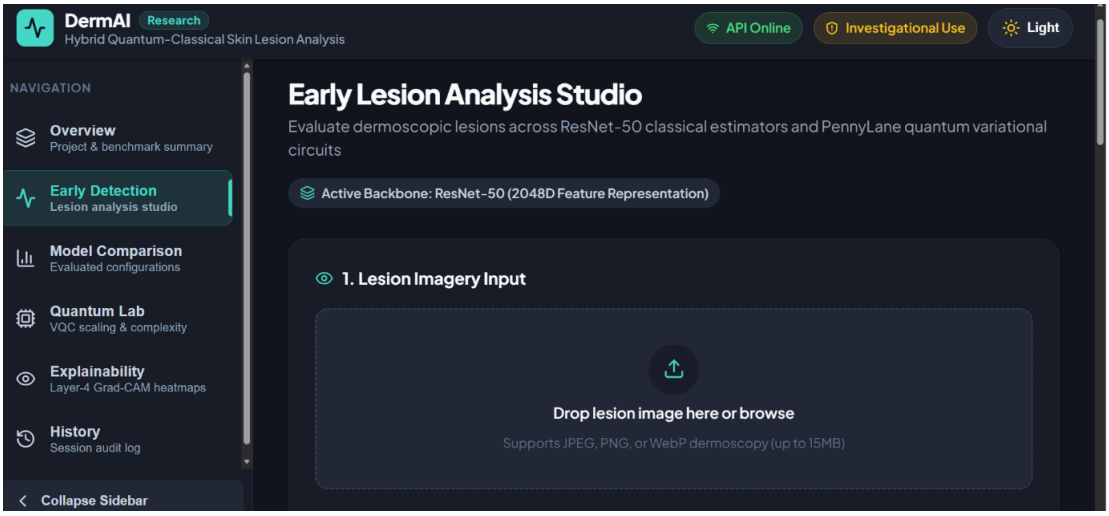


Figure 3.2:Skin Lesion Image Upload

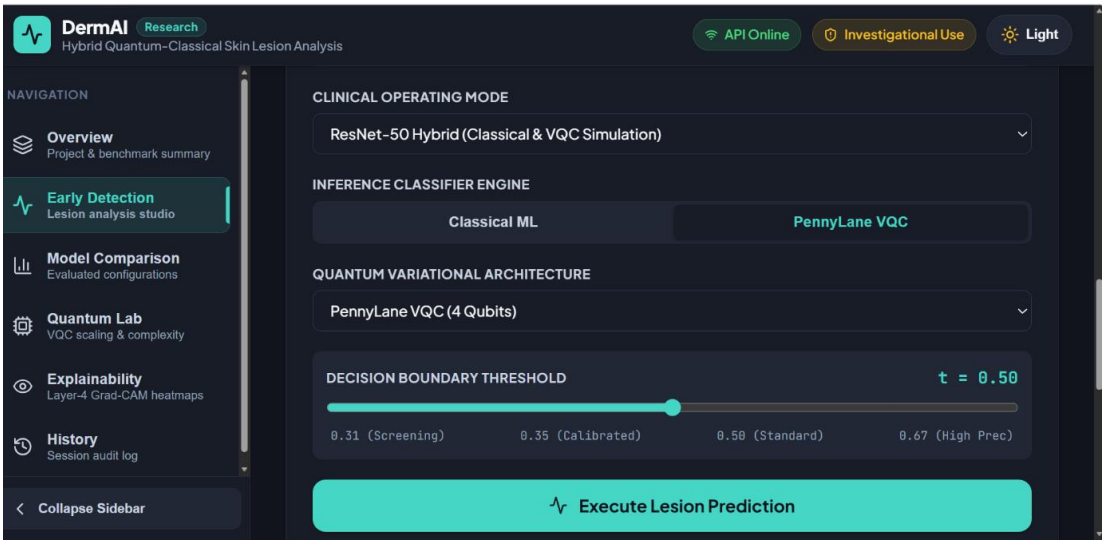


Figure 3.3: Selecting Model for Prediction

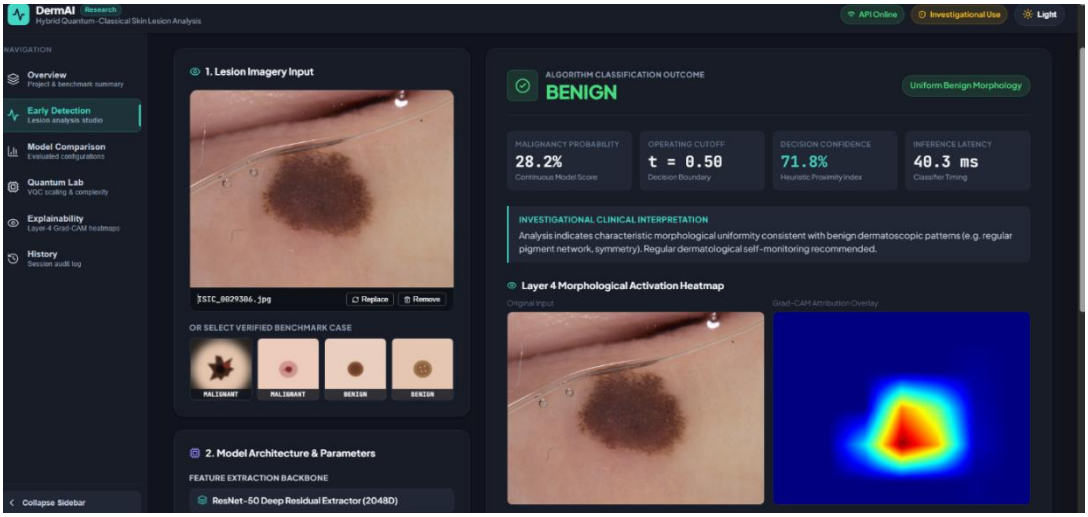


Figure 3.4: Prediction Output Generated by Proposed Model

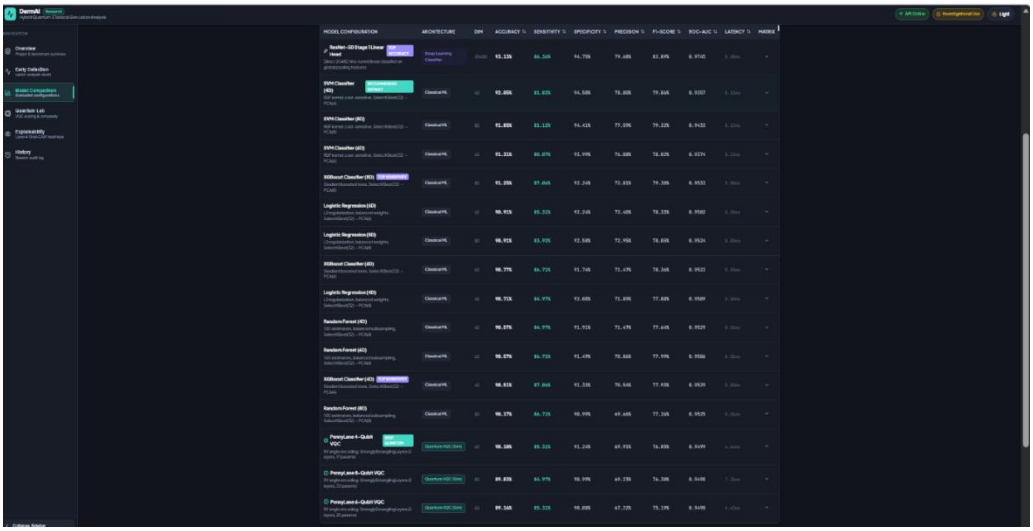


Figure 3.5: Models Comparison

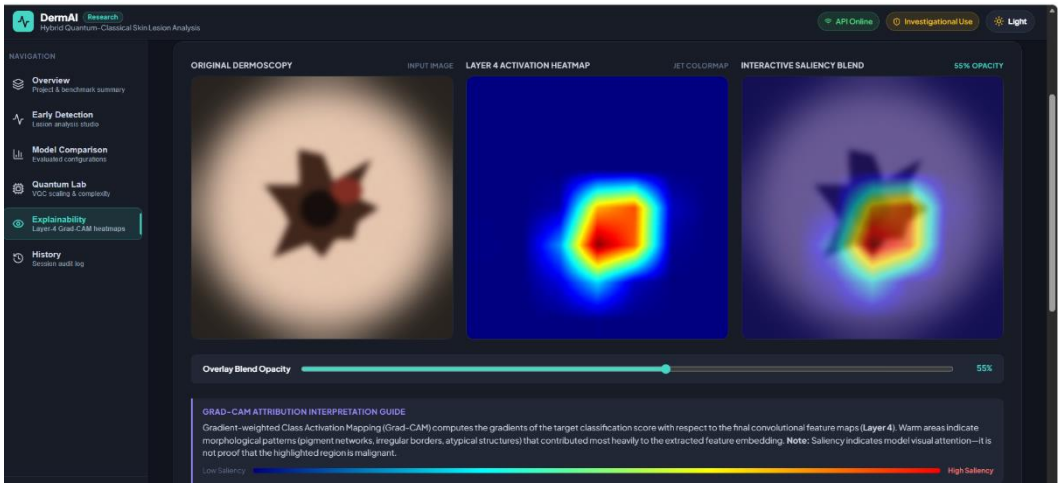


Figure 3.6: Grad cam Based Visual Explaination

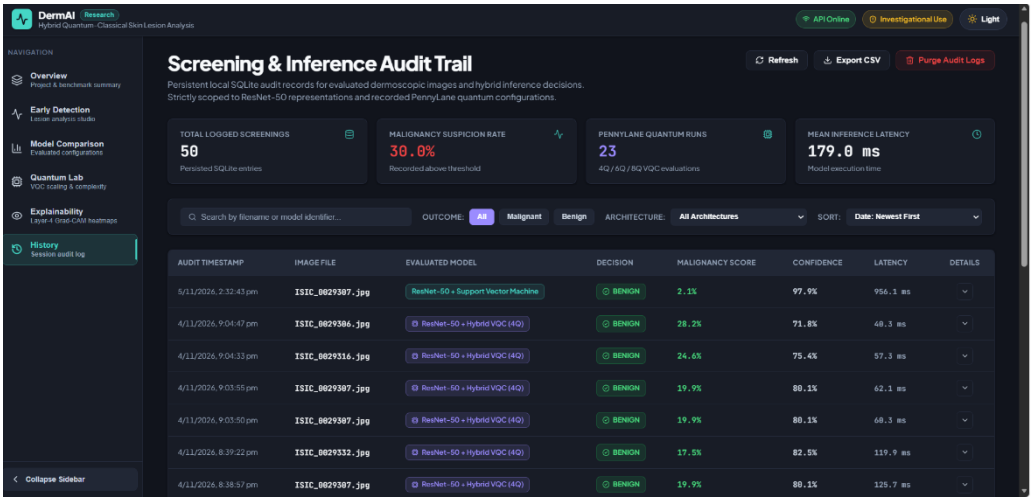


Figure 3.7: Prediction History and Audit Trail

Results and Discussion

The experimental evaluation was conducted on a common held-out test set of 1,485 images, consisting of 1,199 benign and 286 malignant samples. The results demonstrate that the ResNet-50 standalone classification model achieved the highest overall accuracy of 93.13%, with an F1-score of 82.89% and ROC-AUC of 0.9745.

Model	Best Configuration	Accuracy (%)	Precision (%)	Sensitivity (%)	Specificity (%)	F1-score (%)	ROC-AUC
ResNet-50	Direct classification	93.13	79.68	86.36	94.75	82.89	0.9745
SVM	4D	92.05	78.00	81.82	94.50	79.86	0.9357
Logistic Regression	6D	90.91	72.40	85.31	92.24	78.33	0.9502
Random Forest	6D	90.57	70.86	86.71	91.49	77.99	0.9506
XGBoost	8D	91.25	72.81	87.06	92.24	79.30	0.9533
VQC	4 Qubits	90.10	69.91	85.31	91.24	76.85	0.9499

Table 1. Performance Comparison of the Best-Performing Classical and Quantum Models

The results indicate that ResNet-50 achieved the highest accuracy of 93.13% among the evaluated models. Among the compact-feature classical

classifiers, SVM with 4D features achieved the highest accuracy of 92.05%. XGBoost with 8D features obtained the highest sensitivity of 87.06% among the selected configurations, indicating its ability to identify malignant cases. The VQC with 4 qubits achieved an accuracy of 90.10% and sensitivity of 85.31%, showing competitive performance with the compact-feature classical models.

PCA retained 81.56%, 86.91%, and 89.94% of the cumulative explained variance for 4D, 6D, and 8D representations, respectively. This demonstrates the ability to compress high-dimensional ResNet-50 features into compact representations suitable for quantum classification.

Overall, the classical models achieved better accuracy and computational efficiency than the evaluated VQC configurations. However, VQC demonstrated competitive classification performance using a compact feature representation. The framework integrates deep feature extraction, dimensionality reduction, classical and quantum classification, and Grad-CAM-based visual explainability.

CONCLUSION

The proposed explainable hybrid quantum-classical framework demonstrates the potential of combining deep learning, classical machine learning, and quantum machine learning for skin cancer detection. The experimental results show that the ResNet-50 deep learning model achieved the highest accuracy of 93.13%, while classical classifiers such as SVM and XGBoost also demonstrated strong performance using PCA-reduced features. The Variational Quantum Classifier (VQC) achieved an accuracy of 90.10% with four qubits, demonstrating its capability to perform classification using compact feature representations. These findings highlight the effectiveness of deep learning and classical machine learning while also establishing VQC as a promising approach for further exploration in medical image classification. Additionally, the integration of Grad-CAM provides visual explanations of model predictions, improving the interpretability of the proposed framework. The complete system is implemented as a web application, enabling image-based predictions, model comparison, and prediction history management.

Future work will focus on enhancing the performance, scalability, and practical applicability of the proposed framework. The quantum component can be further investigated using advanced quantum feature encoding techniques, optimized variational circuit architectures, improved entanglement strategies, and noise-aware training methods. Evaluating the VQC on real quantum hardware, in addition to simulator-based experiments, would provide further insight into its practical performance and resource requirements. The framework can also be extended using larger and more diverse dermoscopic datasets, multi-class skin lesion classification, external clinical validation, and advanced explainability techniques. Furthermore, optimizing the deep learning and classical classification stages, along with developing more effective hybrid quantum-classical architectures, may improve prediction reliability and computational efficiency. These enhancements could help establish a more robust and interpretable framework for supporting early skin cancer screening and future medical AI research.

REFERENCES

1. Al Mahmud, A., Azam, S., Khan, I. U., Montaha, S., Karim, A., Haque, A., Hasan, M. Z., Brady, M., Biswas, R., & Jonkman, M. (2024). SkinNet-14: A deep learning framework for accurate skin cancer classification using low-resolution dermoscopy images with optimized training time. *Neural Computing and Applications, 36*, 18935–18959. https://doi.org/10.1007/s00521-024-10225-y
2. Aravinda, C. V., Joseph, E. R., & Alasmari, S. (2025). A hybrid quantum-classical approach for multi-class skin disease classification using a 4-qubit model. *IEEE Access, 13*, 114460–114470. https://doi.org/10.1109/ACCESS.2025.3581030
3. Aravinda, C. V., Joseph, E. R., & Alasmari, S. (2025). Optimized DenseNet121 and quantum PennyLane fusion for explainable skin disease recognition and classification. *IEEE Access, 13*, 167957–167969. https://doi.org/10.1109/ACCESS.2025.3608217
4. Farhatullah, X., Chen, X., Zeng, D., Xu, J., Nawaz, R., & Ullah, R. (2024). Classification of skin lesion with features extraction using quantum Chebyshev polynomials and autoencoder from wavelet-transformed images. *IEEE Access, 12*, 193923–193936. https://doi.org/10.1109/ACCESS.2024.3502513
5. Frasca, M., Cutica, I., Pravettoni, G., & La Torre, D. (2025). Optimizing melanoma diagnosis: A hybrid deep learning and quantum computing approach for enhanced lesion classification. *Intelligence-Based Medicine, 12*, Article 100264. https://doi.org/10.1016/j.ibmed.2025.100264
6. Gallazzi, M., Biavaschi, S., Bulgheroni, A., Gatti, T. M., Corchs, S., & Gallo, I. (2024). A large dataset to enhance skin cancer classification with transformer-based deep neural networks. *IEEE Access, 12*, 109544–109559. https://doi.org/10.1109/ACCESS.2024.3439365
7. Gururaj, H. L., Manju, N., Nagarjun, A., Manjunath Aradhya, V. N., & Flammini, F. (2023). DeepSkin: A deep learning approach for skin cancer classification. *IEEE Access, 11*, 50205–50214. https://doi.org/10.1109/ACCESS.2023.3274848
8. Houssein, E. H., Abdelkareem, D. A., Hu, G., Hameed, M. A., Ibrahim, I. A., & Younan, M. (2024). An effective multiclass skin cancer classification approach based on deep convolutional neural network. *Cluster Computing, 27*, 12799–12819. https://doi.org/10.1007/s10586-024-04540-1
9. Hussein, A. A., Montaser, A. M., & Elsayed, H. A. (2025). Skin cancer image classification using hybrid quantum deep learning model with BiLSTM and MobileNetV2. *Quantum Machine Intelligence, 7*, Article 66. https://doi.org/10.1007/s42484-025-00288-y
10. Paraddy, S., & Virupakshappa. (2025). Addressing challenges in skin cancer diagnosis: A convolutional Swin Transformer approach. *Journal of Imaging Informatics in Medicine, 38*(3), 1755–1775. https://doi.org/10.1007/s10278-024-01290-9
11. Reka, S. S., Karthikeyan, H. L., Shakil, A. J., Venugopal, P., & Muniraj, M. (2024). Exploring quantum machine learning for enhanced skin lesion classification: A comparative study of implementation methods. *IEEE Access, 12*, 104568–104584. https://doi.org/10.1109/ACCESS.2024.3434681
12. Suryanarayana, G., Ansari, M. D., K., L. N. C. P., Biksham, V., Swapna, N., & Gogineni, S. (2026). Quantum classical convolutional neural network using ResNet-18 and feature optimization for multiclass skin lesion classification. *Discover Artificial Intelligence, 6*, Article 710. https://doi.org/10.1007/s44163-026-01384-z

13. Tabibi, S. T., Nikravanshalmani, A., & Saboohi, H. (2024). An ensemble classifier based on diverse convolutional neural networks for skin lesions classification. *IEEE Access, 13*, 195673–195686. https://doi.org/10.1109/ACCESS.2024.3442827

14. Toprak, A. N., & Aruk, I. (2024). A hybrid convolutional neural network model for the classification of multi-class skin cancer. *International Journal of Imaging Systems and Technology, 34*(5), Article e23180. https://doi.org/10.1002/ima.23180

15. Yang, G., Luo, S., & Greer, P. (2023). A novel vision transformer model for skin cancer classification. *Neural Processing Letters, 55*, 9335–9351. https://doi.org/10.1007/s11063-023-11204-5