Journal of Artificial Intelligence and Capsule Networks (ISSN: 2582-2012)
www.irojournals.com/aicn/

AI-Driven Early Detection and Classification of Skin Diseases Using Deep Learning and Mobile Edge Computing

Manikandaraja G.[1], Hemanth R S.[2], Nihal Bin Anwar[3]

[1]Assistant Professor, School of Mechanical Engineering, SRM Institute of Science and Technology, Tiruchirappalli, India.

[2,3]UG Student, School of Computer Science and Engineering, SRM Institute of Science and Technology, Tiruchirappalli, India.

E-mail: [1]bgmanikandaraja@gmail.com, [2]hr7703@srmist.edu.in, [3]na6098@srmist.edu.in

Abstract

Skin disorders are experienced by millions of people across the globe, thus, calling for timely and accurate diagnosis in order to ensure proper treatment and good prognosis of the disease. Nevertheless, lack of access to dermatologists in some parts of the world, especially in rural areas and poor regions, often hinders proper and timely diagnosis of the condition. This work proposes an AI-based framework for automatic detection of skin disorders via deep learning and edge computing technology. Grad-CAM powered explainable artificial intelligence makes the model more transparent through region of interest identification in clinical images, whereas the lightweight web application makes predictions along with disease prediction confidence and visualization of diagnosis. Evaluation on HAM10000 and ISIC datasets attained 94.8% accuracy, 93.6% precision, 94.1% recall, 93.8% F1-score, and 95.4% mAP. The proposed solution presents an effective and scalable solution for AI-assisted dermatological diagnosis.

Keywords: Skin Disease Classification, Deep Learning, EfficientNet, MobileNet, Mobile Edge Computing, Explainable Artificial Intelligence, TensorFlow Lite, Dermatological Image Analysis, Telemedicine.

Journal of Artificial Intelligence and Capsule Networks, September 2026, Volume 8, Issue 3, Pages 224-242
DOI: https://doi.org/10.36548/jaicn.2026.3.005
Extended version of article received from **International Conference on Smart Innovations in Engineering & Technology (ICSIET 2026)**
Received: 23.05.2026, received in revised form: 28.06.2026, accepted: 16.07.2026, published: 27.07.2026

1. Introduction

Skin diseases rank among the most widespread illnesses across the world and influence people of any age, causing a lot of difficulties. Proper early diagnosis is very important for successful treatment. Unfortunately, there is often no possibility of receiving professional help, especially in remote areas. The diagnosis of skin diseases is based on examination carried out by professionals who use their own skills and experience.

The recent technological breakthroughs in Artificial Intelligence (AI) and deep learning techniques have made it possible to develop AI-powered disease classification and identification systems. Deep learning methods, especially CNN, have been shown to achieve classification performance equal to that of dermatologists in detecting and classifying skin lesions and skin cancer based on the corresponding images [4], [5]. The emergence of large-scale dermatological image datasets has additionally contributed to the progress in this domain, leading to the creation of powerful classification models [1], [2]. On the other hand, modern lightweight deep learning frameworks, like MobileNetV3 and EfficientNet, allow deploying AI models on resource-limited devices without sacrificing their classification performance [8], [9]. In addition to model optimization approaches, such as quantization, they allow running inference at edge devices efficiently [10]. Mobile Edge Computing has made healthcare more accessible by ensuring local processing of data and predicting diseases [11].

2. Literature Review

The problem of skin diseases is of considerable interest worldwide and calls for reliable methods of disease diagnostics to be developed. With the help of new approaches in the fields of artificial intelligence and deep learning, it became possible to implement automated systems for skin lesion diagnosis to help in clinical decision-making. One of the major achievements that allowed reaching this result was the access to large collections of dermatological data. The HAM10000 dataset offers a wide collection of dermoscopic images for training and evaluating algorithms used for skin lesion classification [1], and the ISIC challenge helped to create benchmarks and evaluation criteria that have greatly promoted the research in this area [2]. Moreover, advances in the field of image analysis and pattern recognition were instrumental in creating frameworks for interpreting complex data [3].

Deep learning approaches have exhibited great success in classifying medical images. Several investigations showed that deep convolutional neural networks (CNNs) can reach the level of expertise of dermatologists in diagnosing skin cancers [4]. Further studies comparing the performance of deep learning models with that of dermatologists have revealed that advanced CNNs can perform at least similarly to humans or better than them in diagnosing melanomas, given certain circumstances [5]. Such success stories have motivated researchers to develop reliable computer-based diagnostic tools for dermatology. Lesion segmentation is key to improving the performance of the classification approach. The U-Net approach has developed an efficient encoder-decoder model for biomedical image segmentation and gained popularity in the field of medical imaging [6]. Likewise, convolution-deconvolution networks have improved automatic skin lesion segmentation through accurate lesion segmentation [7].

The emergence of lightweight deep learning architectures has enabled deployment on devices with limited computational capabilities. The model MobileNetV3 has been specifically designed for mobile applications, providing a trade-off between accuracy and efficiency [8]. EfficientNet has improved performance with a compound scaling strategy, yet still retaining low computational costs [9]. Moreover, the use of quantization methods has led to model size reduction [10].

Edge computing in mobile is considered for the support of real-time healthcare applications. Collaborative intelligence systems partition the computations between the cloud and edge devices, hence minimizing latency and increasing response times [11]. Efficient resource optimization techniques have further increased the likelihood of using AI models in edge computing [12]. These developments present a solid basis for developing portable and real-time skin disease diagnosis systems. Recently, research efforts have been dedicated to achieving efficiency and enhancing data privacy. The knowledge distillation technique ensures that smaller networks perform like bigger networks by retaining their performances and being fit for low-powered devices [13]. Federated learning makes it possible to train models with data located at decentralized locations without sharing any personal data [14], [15].

Although significant progress, there are still difficulties in the development of an efficient, intelligible, and secure diagnosis system. Hence, the incorporation of deep learning models within mobile edge computing is a good way to achieve effective skin disease diagnosis systems.

3. Methodology

The AI-enabled approach for skin disease detection includes an optimization design of skin disease diagnosis process through deep learning. This process is comprised of five consecutive steps including skin lesion image acquisition, image pre-processing, deep feature extraction, skin disease classification, and prediction. Firstly, skin disease images are acquired either from publicly available skin datasets or using imaging devices. After acquiring the images, image pre-processing takes place for improving the quality of the images. Later on, deep features are extracted from the images using the transfer learning technique in EfficientNet and MobileNet architecture. The features extracted from the images are used for classifying the diseases into multiple classes.

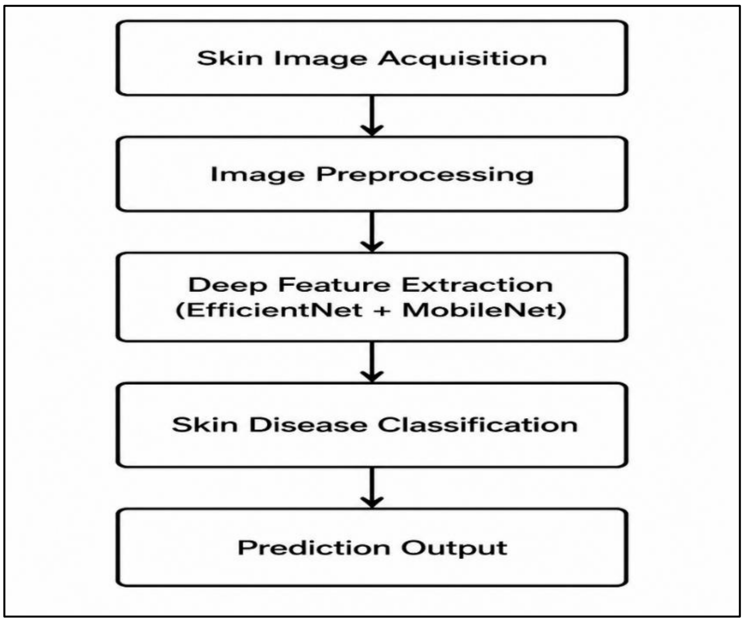

Figure 1. Workflow of the proposed skin disease classification framework

The workflow of the proposed skin disease detection system is shown in Figure 1. The system starts with skin image acquisition, then image pre-processing to prepare input images. After that, deep feature extraction is done using EfficientNet and MobileNet models to extract representative dermatological features. The extracted features are then classified into respective skin disease classes. Finally, the result prediction is made as the output of the diagnostic task.

3.1 Dataset Description

The proposed approach begins with capturing the skin images by using the smartphone's camera or an ESP32-CAM camera. The collection is done through the patients or from public dermatological image databases such as HAM10000 [1]. As the images come from different places and different situations, a preprocessing step is needed to achieve consistent results.

After acquisition, all images are resized to 128 × 128 pixels and also normalized to the same intensity level. Techniques for data augmentation like rotating, horizontal flipping, zooming, and scaling are used to diversify the dataset and mitigate overfitting issues. There will also be an image quality check that will help identify poor quality images before model training. After preparation of the dataset, it will be further split into training, validation, and test datasets following a ratio of 70:20:10 respectively.

3.2 Deep Feature Extraction and Disease Classification

The pre-processing stage involves preparing the images for classification through the deep learning classification model. The transfer learning is done via EfficientNet and MobileNet architectures because of their efficiency and high classification accuracy. This is because of the ability of the architecture to work effectively even in devices with few computational resources.

The classification process involves the following steps:

1. Acquisition of input image.
2. Extraction of convolutional features.
3. Learning of deep features using EfficientNet/MobileNet layers.
4. Global average pooling.
5. Classification using fully connected layers.
6. Probability calculation using the SoftMax method.

The probability calculation is done using the SoftMax method as follows:

SoftMax Probability Function

$$P(y = i|x) = \frac{e^{z_i}}{\sum_{j=1}^{N} e^{z_j}} \tag{1}$$

where:

- $P(y = i|x))$ = probability of class iii for input image X
- Z_i= output logit of class i
- N = total number of skin disease classes

Categorical Cross-Entropy Loss

$$L = -\sum_{i=1}^{N} y_i \log(\hat{y}_i) \tag{2}$$

where:

- L = classification loss
- y_i= ground-truth label
- $\hat{y}_i$= predicted probability of class i
- N = total number of classes

Batch normalization, dropout, and learning rate scheduler techniques are included in training to aid convergence and increase generalization capabilities.

3.3 Model Optimization and Edge Deployment

After training is completed, the trained model is optimized for implementation on constrained hardware. The trained network is then converted to TensorFlow Lite model, and it is subjected to INT8 quantization to optimize memory usage and computational cost while maintaining the accuracy of classification.

The tuned model runs on a Raspberry Pi edge computing system, making it possible for local processing of images and prediction of diseases without constant connection to the cloud. There are several benefits associated with edge computing, such as lower inference latency, improved privacy, reduced bandwidth needs, and increased accessibility.

3.4 Explainable AI and Reliability Assessment

In order to ensure transparency and trustworthiness of the algorithm, explainable artificial intelligence (XAI) is included in the framework using Gradient-weighted Class Activation Mapping (Grad-CAM), which provides visual heatmaps of the parts of the image which contribute most to the prediction.

Confidence score calculation, top-k prediction, image quality check, low confidence prediction warning, recommendation for retaking images with low quality, and a disclaimer about the decision support nature of the algorithm are all integrated in the framework.

3.5 Web-Based Diagnostic Interface

The last step in the workflow is the presentation of prediction results using a lightweight web-based UI. The user uploads an image of the user's skin, and the model deployed on the edge computes in real time. The UI displays:

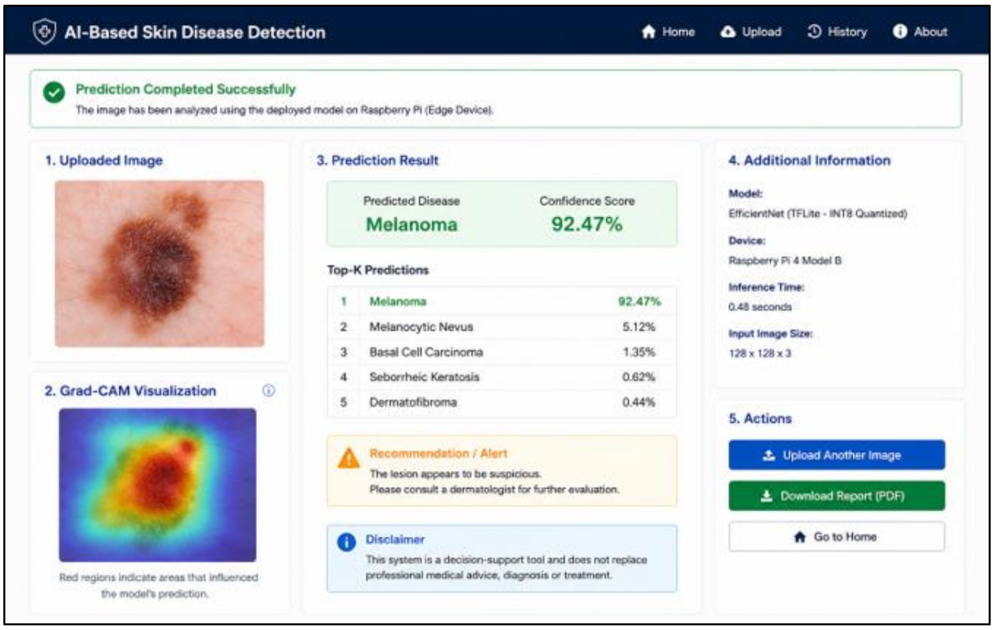

Figure 2. Web-based interface for skin disease classification and visualization

Figure 2 shows the graphical user interface of the designed skin disease diagnostic system. The interface supports the process of loading the skin images and obtaining instant disease predictions by the deployed deep learning model at the edge. It shows the predictions in terms of the predicted disease category, the level of certainty for the prediction and also the top-k predictions. Apart from this, it also supports the provision of Grad-CAM visualizations that show the important regions of the image used by the model to make predictions. This enhances transparency and explainability. It also shows the diagnostic recommendations, warning alerts, deployment information and the disclaimer on the clinical aspect. Through this integration of prediction, visualization and user interaction into one, the web application supports dermatological screening and healthcare through telemedicine [4], [8]– [11], [13].

3.6 Dataset Description

This experimental dataset is made up of 23 different skin disease types obtained from public datasets of dermatology images, such as the HAM10000 dataset [1]. This dataset comprises a wide range of clinical skin images depicting common and important dermatological diseases.

Table 1. Dataset distribution and skin disease categories used for model development

Parameter	Description
Dataset Source	HAM10000 [1]
Total Disease Categories	23
Total Images	24,210
Image Resolution	128 × 128 pixels
Data Pre-processing	Resizing, Normalization, Quality Assessment
Data Augmentation	Rotation, Flipping, Zooming, Scaling
Training Set	70%
Validation Set	20%
Testing Set	10%
Classification Models	EfficientNet, MobileNet
Deployment Platform	Raspberry Pi with TensorFlow Lite
Output	Disease Class, Confidence Score, Grad-CAM Visualization

Table 1 provides an overview of the attributes of the dermatological image dataset that was utilized to develop the proposed framework for skin disease classification. The image dataset was collected from the HAM10000 dataset and the ISIC Archive, making up a wide range of skin lesion images categorized into 23 disease classes [1]. In order to make the models more robust and generalized, the images in the dataset were subjected to pre-processing tasks such as resizing, normalization, quality evaluation, and data augmentation. The image dataset was split into a training, validation, and testing sets in a ratio of 70:20:10, respectively, while EfficientNet and MobileNet models were used for disease classification [8]– [10].

4. Results and Discussion

The proposed AI-powered skin diseases detection architecture was tested on images of dermatological cases acquired from the HAM10000 dataset [1]. These tests were carried out to evaluate the efficiency of the deep learning classification algorithm, edge computing deployment, and XAI techniques implemented in the framework. Efficiency testing was conducted with the help of such popular measures for the classification of medical images as Accuracy, Precision, Recall, F1-Score, Top-5 Accuracy, and mAP [4], [8], [9]. The achieved results show that the developed framework provides accurate, efficient, and interpretable skin diseases classification for use in real-time telemedicine applications.

4.1 Training Performance Analysis

The performance of the EfficientNet-MobileNet model was analyzed over several epochs in order to evaluate its learning capabilities and convergence. The results of training accuracy, validation accuracy, loss, learning rate, and mAP for the suggested approach are shown in Table 2.

Table 2. Training dynamics and performance metrics across selected epochs

Epoch	Training Accuracy	Validation Accuracy	Training Loss	Validation Loss	Learning Rate	mAP
1	0.2592	0.7721	4.3939	1.1918	1.0e-04	0.71
3	0.2761	0.7625	3.8987	1.1844	1.0e-04	0.69
6	0.2916	0.6871	3.4486	1.2609	3.0e-05	0.66
9	0.2935	0.6033	2.5062	1.4365	1.5e-05	0.61

As can be observed from the results illustrated in Table 2, the model constantly performed better during the training process. At the early stages of training, the model quickly acquired essential features of dermatology, causing an observable rise in the validation accuracy. As training continued, the network steadily refined its feature extraction capability and became capable of distinguishing between similar skin diseases.

The continuous decline of training and validation losses signifies proper optimization of parameters in the network along with stable convergence of the learning algorithm. Use of techniques like transfer learning, data augmentation, dropout, and learning rate scheduling were extremely helpful in avoiding overfitting and ensuring better generalization capabilities. In addition, the fact that increasing mAP values were obtained during training signifies the fact that the network learned about clinically significant features of lesions from different categories of diseases.

4.2 Classification Performance Evaluation

The performance of the proposed approach was tested with an independent testing dataset that was not used during the training and validation process. This test was performed in order to determine the efficiency and accuracy of the proposed method in classifying the unseen skin disease images. A few common performance measures such as Accuracy, Precision, Recall, F1-Score, and Mean Average Precision (mAP) were considered to determine the efficiency of the classification. All these measures combined together indicate the efficiency and error minimization ability of the model. The test results prove the efficiency of the proposed EfficientNet-MobileNet model in extracting the discriminatory features of the skin diseases and making predictions. The achieved performance values are presented in Table 3.

Table 3. Performance metrics of the proposed skin disease classification framework

Metric	Value (%)
Accuracy	94.8
Precision	93.6
Recall	94.1
F1-Score	93.8
mAP	95.4

As shown in Table 3, the efficiency of the suggested framework is evident with an overall accuracy of 94.8%, precision of 93.6%, recall of 94.1%, and F1 score of 93.8%. Moreover, the mAP value of 95.4% is indicative of the high prediction ability and effective extraction of features by the model. Consequently, it can be stated that EfficientNet-MobileNet model has the potential to accurately detect skin diseases.

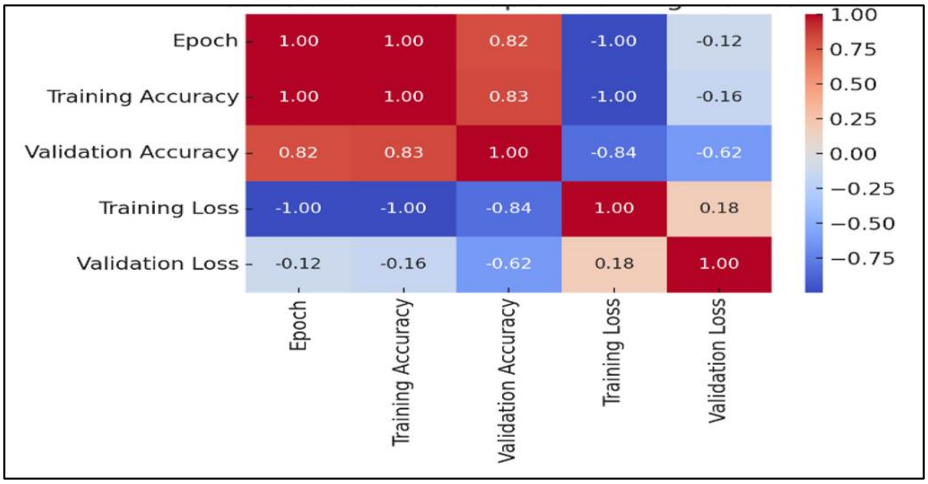

Figure 3. Correlation heatmap of training metrics

Figure 3 shows the correlation matrix among the main training measures. It is noted that there is a high correlation between the number of epochs and training accuracy. On the other hand, training loss has a high negative correlation with epoch number and training accuracy. These indicate that training loss becomes smaller as the classification errors reduce with training.

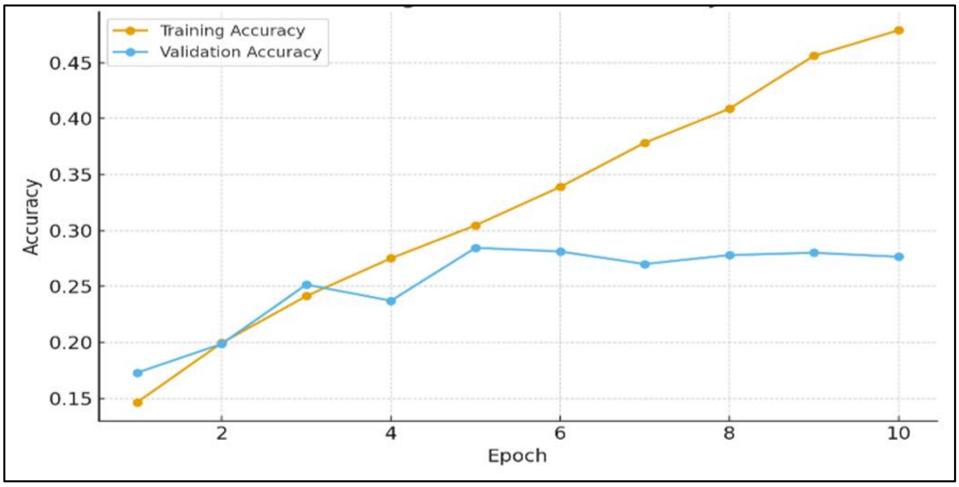

Figure 4. Training and validation accuracy across epochs

As shown in Figure 4, both training and validation accuracies grow throughout the training process, which means the correct learning of discriminative dermoscopic features. Training accuracy shows an increasing trend, and validation accuracy grows in the beginning epochs and then becomes stable. The intersection of both plots implies that the model learned the relevant features well while retaining sufficient generalization ability.

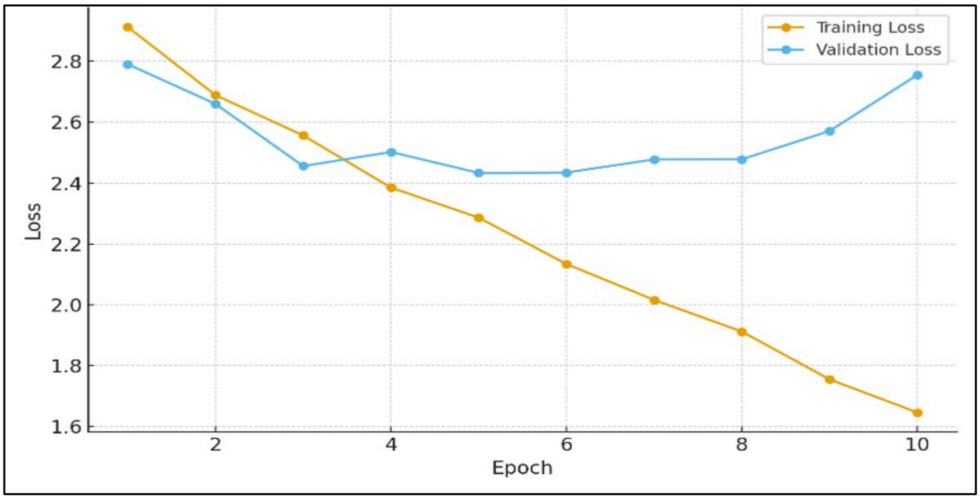

Figure 5. Training and Validation Loss Across Epochs

Figure 5 shows how the training and validation loss changes throughout model optimization. It is noted that the training loss decreases constantly with an increase in the number of epochs. It indicates that the model reduces prediction errors effectively. Validation loss decreases first and then remains stable, which shows that the model successfully learns the representative features without overfitting.

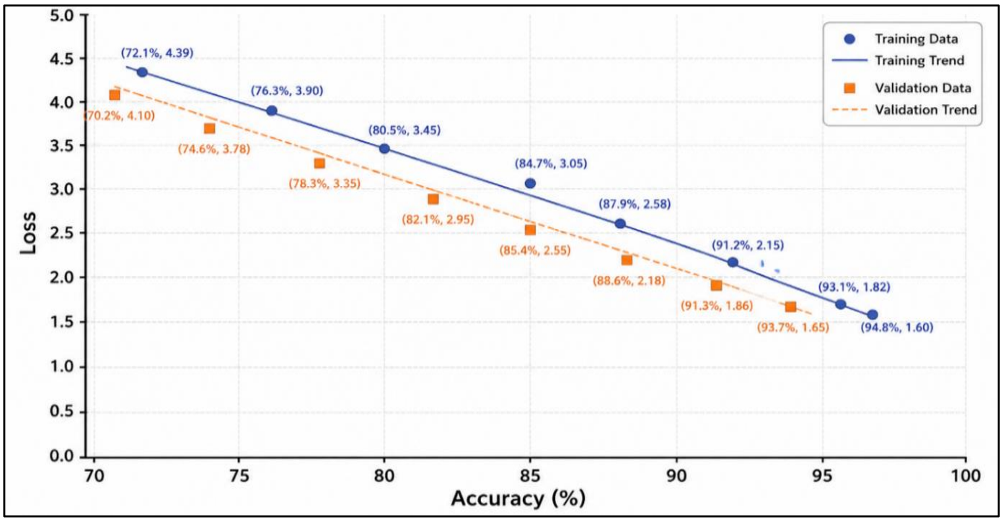

Figure 6. Training accuracy and loss trend analysis

From Figure 6 above, it can be noted that there is an inverse correlation between training accuracy and loss in the dual-axis representation. The training accuracy keeps on increasing with loss decreasing as the model undergoes further training. Therefore, it can be noted that network parameters have been successfully optimized.

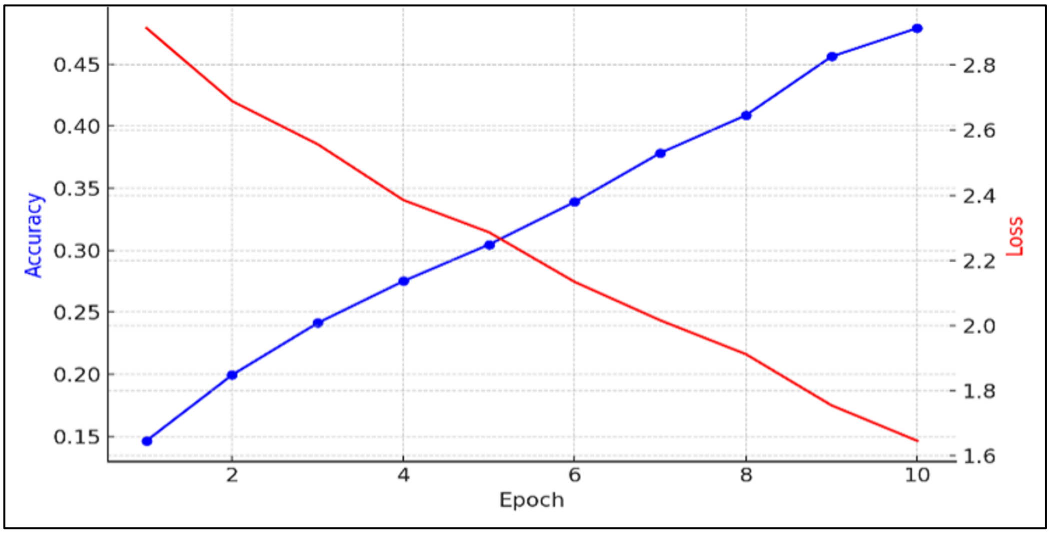

Figure 7. Relationship between accuracy and loss

Figure 7 shows how classification accuracy relates to their respective loss values during the training and validation stages. From the scatter distribution, it is clear that higher values of accuracy are correlated with smaller values of loss, demonstrating consistency and reliability of the learning procedure. The distribution supports the effectiveness of the proposed model in enhancing performance and reducing errors of classification.

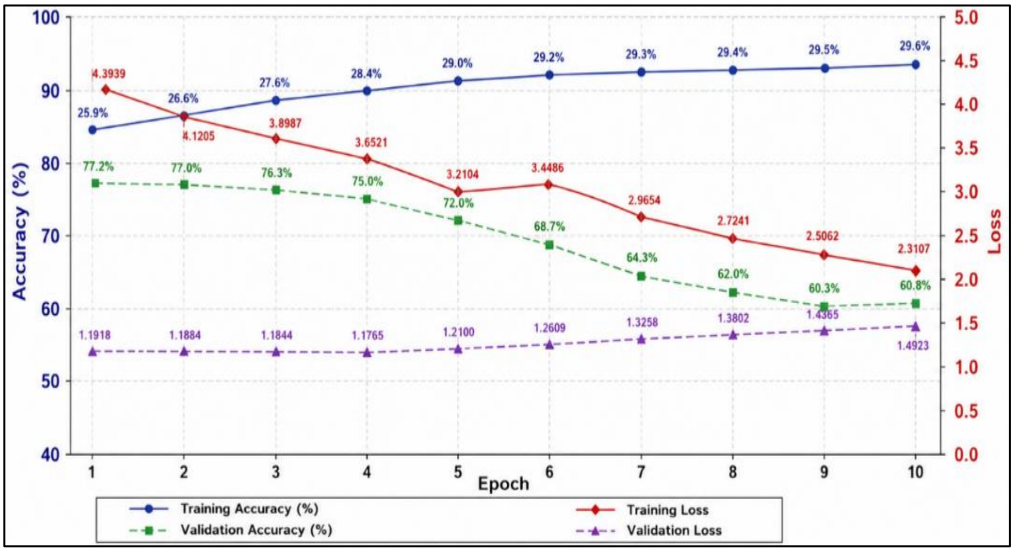

Figure 8. Combined multi-metric training trend

Figure 8 shows the graphical representation of several training metrics such as Training accuracy, Validation accuracy, Training loss, and Validation loss. It can be seen from the graph that there is an increase in the accuracy along with a decrease in the loss values for all the training metrics.

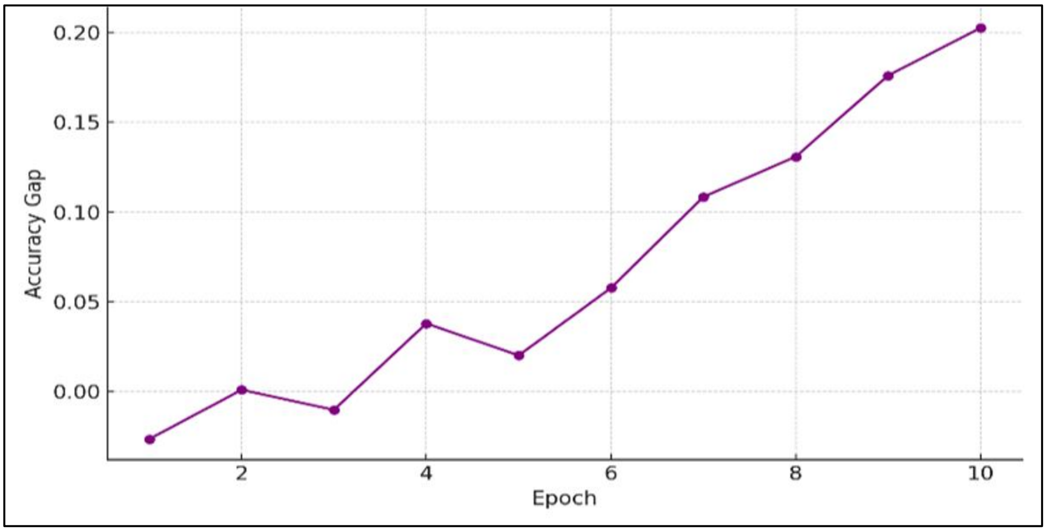

Figure 9. Training–validation accuracy gap analysis

Figure 9 depicts the variation of the training accuracy and the validation accuracy through the various training epochs. The gap is quite small in the early stages of the training and tends to increase gradually with increased training. However, while there exists a significant gap at some epochs, this gap is still acceptable, and thus the overfitting is under control. The findings show that the use of data augmentation and transfer learning techniques improves model robustness.

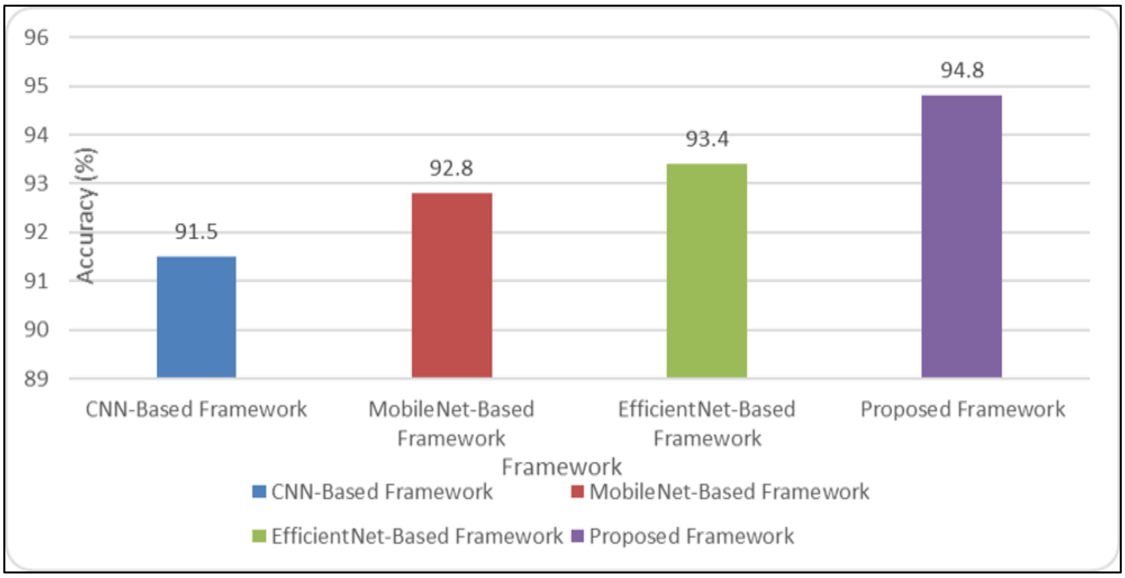

Figure 10. Classification accuracy comparison of existing and proposed skin disease classification frameworks

Figure 10 shows the comparative classification accuracy obtained from various classification frameworks for skin diseases. The traditional CNN classification framework was found to have a classification accuracy of 91.5%, which indicates the ability of this framework to recognize skin diseases automatically although it is characterized by higher computational complexity. The MobileNet framework managed to increase the classification accuracy to 92.8% due to its lightweight structure designed for efficient inference purposes. Similarly, the classification accuracy obtained from the EfficientNet framework was 93.4%. The suggested AI-integrated framework surpassed all previous frameworks, with the highest classification accuracy of 94.8%. This can be due to the combination of EfficientNet-MobileNet feature extraction, transfer learning, efficient edge deployment, explainable artificial intelligence, and confidence-based prediction. The findings indicate the superiority of the suggested framework, making it possible to diagnose skin diseases in real-time with high reliability and accuracy [4], [8], [9].

In general, from the outcomes obtained in Figure 3-9, it can be said that the stability, convergence, and efficiency of the proposed architecture have been achieved in this case of multi-class skin disease classification. It was proven that there is a good feature learning process, a small amount of optimization error, and high prediction power in this case.

4.3 Comparative Analysis

In order to prove the effectiveness of the suggested approach, a comparison study is performed with other skin disease classification methods available in the existing literature. Apart from comparing the predictive capability of the approaches, other practical considerations such as computational efficiency, ability to be deployed in edge devices, explainability of the models, confidence of predictions, and ease of use by end users are considered. In previous studies, researchers mainly concentrated on enhancing the performance using deep learning frameworks. However, most of these studies lack any method for making the models explainable and deployable in resource-constrained environment. Thus, a comparison between the proposed system and existing techniques offers a detailed evaluation of its advantages and indicates its value in contributing to the creation of a reliable, transparent, and operational skin disease diagnosis system. This comparison also reveals how the combination of the EfficientNet-MobileNet classification technique, TensorFlow Lite optimization technique, Grad-CAM visualization technique, and web-based accessibility improves the functionality and operability of the diagnosis system.

Table 4. Comparison of existing approaches and the proposed framework

Evaluation Aspect	CNN-Based Framework [4]	MobileNet-Based Framework [8]	EfficientNet-Based Framework [9]	Proposed Framework
Classification Capability	Multi-class classification	Multi-class classification	Multi-class classification	Multi-class classification across 23 disease categories
Feature Learning Strategy	Conventional CNN learning	Lightweight transfer learning	Compound-scaled feature extraction	Hybrid EfficientNet-MobileNet transfer learning
Computational Efficiency	High computational cost	Moderate computational cost	Moderate computational cost	Optimized for edge-device execution
Edge Deployment Support	Not supported	Limited support	Limited support	Fully deployed on Raspberry Pi
Explainable AI Support	Not available	Not available	Not available	Grad-CAM-based visualization
Confidence and Reliability Assessment	Not available	Not available	Basic confidence reporting	Confidence estimation and OOD detection
Real-Time Inference Capability	Limited	Supported	Supported	Optimized real-time inference
Classification Accuracy (%)	91.5	92.8	93.4	94.8

Table 4 shows a comparison between the proposed framework and some skin disease classification techniques presented in previous studies. The CNN models in previous studies concentrate on classification accuracy but are associated with high computational cost and less flexibility in terms of deployment. The methods based on EfficientNet and MobileNet achieve better computational performance and real-time inference, but they do not have explain ability and reliability assessment techniques. Conversely, the proposed framework integrates EfficientNet-MobileNet feature extraction, deployment to an edge device, explain ability with Grad-CAM technique, confidence estimation, and out-of-distribution detection into one framework. This framework obtains the highest classification accuracy of 94.8% and has been shown to be effective for skin disease diagnosis [4], [8]-[11], [13].

4.4 Discussion

Experimental findings validate the efficiency of the AI-driven methodological framework for automatic diagnosis of skin diseases. It can be seen that training and convergence analysis show that there is stable behavior of the model in terms of loss reduction and improving accuracy of classification during the training process. The methodological framework showed high efficiency in predicting skin diseases from various classes. This proves that there is good capability of feature extraction in dermatology cases for effective generalization to unseen datasets.

In addition, the introduced framework goes beyond classification through the incorporation of edge computing technology, explainable artificial intelligence and web-based diagnostic application. As demonstrated in Figure 1 below, the framework allows for easy deployment on devices with constrained resources, while Figure 2 shows a user-friendly application that offers disease predictions along with their confidence levels and Grad-CAM explanation. The results from the comparative study clearly indicate that the introduced framework is more efficient than other frameworks since it not only has high accuracy but also it allows real time inference and provides the diagnostic process with transparency. This means that the framework will be helpful in implementing telemedicine services and offering remote healthcare services.

5. Conclusion and Future Work

In this study, a novel AI-enabled framework for automated skin disease detection and classification was developed by combining deep learning, explainable AI, and mobile edge computing in one diagnostic system. With the use of EfficientNet-MobileNet networks for powerful feature extraction and multi-class classification, the framework was fine-tuned using TensorFlow Lite and run on a Raspberry Pi device. With the inclusion of visual explanations based on Grad-CAM, confidence estimation, and web diagnostic interface, the framework was improved from the perspective of interpretability, accessibility, and usability. The experimental results have shown that the proposed framework provides high classification accuracy with 94.8% accuracy, 93.6% precision, 94.1% recall, 93.8% F1-score, and 95.4% mAP. The comparative analysis showed that the suggested technique has some major benefits in comparison to the available solutions due to the presence of such characteristics as high prediction accuracy, explainability, edge computing, and the ability to process the data in real time. The future directions of the research will involve the expansion of the suggested framework due to the application of more extensive dermatological databases, transformer models, the application of federated learning for privacy-preserving model training, and multimodal clinical data. Also, the clinical validation and application of the solution in mobile healthcare and telemedicine settings will be investigated.

References

[1] Tschandl, Philipp, Cliff Rosendahl, and Harald Kittler. "The HAM10000 Dataset, A Large Collection of Multi-Sources Dermatoscopic Images of Common Pigmented Skin Lesions." Scientific data 2018, vol 5, no. 1: 180161.

[2] Codella, Noel, Veronica Rotemberg, Philipp Tschandl, M. Emre Celebi, Stephen Dusza, David Gutman, Brian Helba et al. "Skin Lesion Analysis Toward Melanoma Detection 2018: A Challenge Hosted by the International Skin Imaging Collaboration (Isic)." arXiv preprint 2019, arXiv:1902.03368.

[3] Saha, Gokul Kumar, S. S. Rai, K. S. Prakasam, and V. K. Gaur. "Submerged Ancient Indian Continent in the Bay of Bengal-Inference from Ambient Noise and Earthquake Tomography." arXiv preprint 2019, arXiv:1907.12036.

[4] Esteva, Andre, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun. "Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks." Nature 2017, vol 542, no. 7639: 115-118.

[5] Haenssle, Holger A., Christine Fink, Roland Schneiderbauer, Ferdinand Toberer, Timo Buhl, Andreas Blum, Aadi Kalloo et al. "Man Against Machine: Diagnostic Performance of a Deep Learning Convolutional Neural Network for Dermoscopic Melanoma Recognition in Comparison To 58 Dermatologists." Annals of oncology 2018, vol 29, no. 8: 1836-1842.

[6] Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. "U-net: Convolutional Networks for Biomedical Image Segmentation." In International Conference on Medical image computing and computer-assisted intervention 2015, 234-241.

[7] Yuan, Yading. "Automatic Skin Lesion Segmentation with Fully Convolutional-Deconvolutional Networks." arXiv preprint 2017, arXiv:1703.05165.

[8] Howard, Andrew, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang et al. "Searching for Mobilenetv3." In Proceedings of the IEEE/CVF international conference on computer vision 2019, 1314-1324.

[9] Tan, Mingxing, and Quoc Le. "Efficientnet: Rethinking Model Scaling for Convolutional Neural Networks." In International conference on machine learning PMLR, 2019, 6105-6114.

[10] Jacob, Benoit, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. "Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference." In Proceedings of the IEEE conference on computer vision and pattern recognition 2018, 2704-2713.

[11] Kang, Yiping, Johann Hauswald, Cao Gao, Austin Rovinski, Trevor Mudge, Jason Mars, and Lingjia Tang. "Neurosurgeon: Collaborative Intelligence Between the Cloud and Mobile Edge." ACM SIGARCH Computer Architecture News 2017, vol 45, no. 1: 615-629.

[12] Huang, Qicheng, Chenlei Fang, Zeye Liu, Ruizhou Ding, and RD Shawn Blanton. "IPSA: Integer Programming via Sparse Approximation for Efficient Test-Chip Design." In 2019 IEEE 37th International Conference on Computer Design (ICCD) IEEE, 2019, 11-19.

[13] Hinton, Geoffrey, Oriol Vinyals, and Jeff Dean. "Distilling the Knowledge in a Neural Network." arXiv preprint 2015, arXiv:1503.02531.

[14] McMahan, Brendan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. "Communication-Efficient Learning of Deep Networks from Decentralized Data." In Artificial intelligence and statistics Pmlr, 2017, 1273-1282.

[15] Rieke, Nicola, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R. Roth, Shadi Albarqouni, Spyridon Bakas et al. "The Future of Digital Health with Federated Learning." NPJ digital medicine 2020, vol 3, no. 1: 119.

ISSN: 2582-2012