

Skin Docnet: Real-Time Melanoma Vs Benign Classification Using Efficientnet

Prof. Ashwini Sangam¹, Sandeep²

¹Assistant Professor, Department of Computer Science and Engineering (MCA)
Visvesvaraya Technological University, Centre for PG Studies, Kalaburagi, Karnataka, India

²Student, Department of Computer Science and Engineering (MCA)
Visvesvaraya Technological University, Centre for PG Studies, Kalaburagi, Karnataka, India

Abstract - Melanoma is one of the most serious forms of skin cancer, and detecting it at an early stage is important for improving treatment outcomes. Conventional diagnosis mainly depends on visual examination by dermatologists, which can be time-consuming and may vary between observers. This study presents Skin DocNet, a web-based deep learning system designed for the real-time classification of skin lesions into melanoma and benign categories using the EfficientNet architecture. The proposed system accepts a skin lesion image from an authenticated user, performs image validation and preprocessing, and then applies the trained EfficientNet model to generate a classification result with a confidence score. To improve interpretability, the system incorporates Grad-CAM, which highlights the image regions that contribute to the model's prediction. The application is developed using Python, TensorFlow/Keras, OpenCV, and Flask, providing an accessible interface for image upload and result visualization. Model performance is assessed using standard classification measures, including accuracy, precision, recall, F1-score, and a confusion matrix. The proposed system is intended as a preliminary screening and decision-support tool that can assist users and healthcare professionals in identifying potentially suspicious skin lesions. It is not intended to replace professional dermatological examination or clinical diagnosis.

Key Words : Melanoma, Benign Skin Lesion, EfficientNet, Deep Learning, Convolutional Neural Network, Grad-CAM, Skin Cancer Detection, Image Classification.

1.INTRODUCTION

Skin cancer is a major dermatological concern in which abnormal growth of skin cells can develop into potentially dangerous lesions. Among the different types of skin cancer, melanoma is particularly concerning because of its ability to spread to other parts of the body when it is not detected and treated at an early stage. Identifying suspicious lesions promptly can therefore play an important role in improving clinical outcomes. However, distinguishing melanoma from benign skin lesions can be difficult because some lesions may have similar visual characteristics.

Conventional skin lesion assessment generally involves clinical examination and dermoscopic inspection by trained dermatologists. This process requires considerable expertise

and can become time-consuming when large numbers of patients need to be examined. Differences in clinical experience may also result in variation when interpreting visually similar lesions. These limitations have encouraged researchers to investigate computer-based methods that can support dermatological image analysis.

Recent developments in artificial intelligence and deep learning have created new possibilities for automated medical image classification. Convolutional neural networks can learn visual characteristics directly from image data instead of depending entirely on manually designed features. Such characteristics may include colour patterns, lesion boundaries, texture variations, and structural features. Deep learning models can therefore be used to assist in distinguishing suspicious lesions from non-malignant ones.

In this work, Skin DocNet is developed as a web-based system for melanoma and benign skin lesion classification using the EfficientNet architecture. EfficientNet is selected because it provides a balance between classification capability and computational efficiency by systematically scaling network depth, width, and input resolution. The trained model receives a processed skin lesion image and produces a predicted category together with a confidence value. The system also incorporates image validation so that unsuitable or irrelevant inputs can be rejected before classification.

To improve the interpretability of the prediction process, the proposed system uses Grad-CAM to generate visual explanations of the model's decision. The resulting heatmap indicates the regions that contributed to the classification, providing additional information about the model's focus during prediction. The application is implemented using Python, TensorFlow/Keras, OpenCV, and Flask, with a web interface for authentication, image submission, prediction, and result presentation. The overall workflow described in the project consists of image upload, validation, preprocessing, EfficientNet-based classification, Grad-CAM generation, and result visualization.

The primary purpose of Skin DocNet is to provide preliminary screening and decision support rather than replace professional medical diagnosis. The system is intended to demonstrate how deep learning and web technologies can be combined to create an accessible tool for supporting the early assessment of skin lesions.

2.PROBLEM STATEMENT

Accurate identification of melanoma from benign skin lesions is a challenging task because different lesions can exhibit similar visual characteristics. Conventional assessment depends largely on dermatological expertise and examination of lesion images, which can require significant time and specialized knowledge. In situations where expert consultation is limited, obtaining an initial assessment may be difficult.

Another challenge is the increasing number of skin-related cases that require examination. Manually evaluating a large collection of lesion images can increase the workload of healthcare professionals and may introduce differences in interpretation. Therefore, there is a need for a computer-assisted approach that can analyze skin lesion images consistently and provide an initial classification.

The problem addressed in this study is the development of an automated system capable of distinguishing melanoma from benign skin lesions using image-based deep learning. The system should accept suitable skin images, perform the required preprocessing, classify the lesion using an EfficientNet-based model, and provide the prediction along with an associated confidence value. It should also prevent unsuitable or irrelevant images from being processed. The proposed system further aims to provide an interpretable visualization of the model's decision using Grad-CAM.

The developed system is intended to function as a preliminary screening and decision-support solution. It should assist users in obtaining an initial assessment of a skin lesion while recognizing that the final diagnosis must be made by a qualified medical professional. The project report similarly defines the system as a decision-support tool rather than a replacement for medical professionals.

3.OBJECTIVES

The main objective of the proposed Skin DocNet system is to develop a deep learning-based web application capable of distinguishing melanoma skin lesions from benign lesions using the EfficientNet architecture. The system is designed to provide an automated approach to skin lesion image analysis, reducing the dependence on manual visual assessment for preliminary screening. The objectives are formulated around the complete process, beginning with image preparation and continuing through classification, interpretation, and result presentation.

The proposed system aims to prepare skin lesion images through appropriate preprocessing so that the input supplied to the EfficientNet model is standardized and suitable for classification. The project also focuses on implementing image validation to ensure that only appropriate skin-related images are processed. This is important because accepting irrelevant images could result in meaningless predictions and reduce the reliability of the application. The original project specifically includes validation mechanisms for uploaded images before they are passed to the prediction stage.

Another objective is to develop an EfficientNet-based classification model that can learn relevant visual characteristics from skin lesion images. The model is intended to identify patterns associated with melanoma and benign lesions by learning characteristics such as colour variation, texture, and lesion structure. EfficientNet is selected because the project considers its balance between classification capability and computational requirements important for an application intended to provide predictions efficiently.

The project also aims to provide a simple and accessible web interface for image-based prediction. The application uses Flask as the web framework and provides functionality for image uploading, prediction, and result visualization. The system architecture includes user interaction, backend processing, image preprocessing, model inference, and presentation of the final result.

An additional objective is to provide meaningful information along with the classification output. Instead of displaying only the predicted class, the system is designed to provide a confidence score and supporting information to make the result easier to understand. Grad-CAM is also incorporated to generate a visual representation of the regions that contribute to the model's prediction. This provides an additional layer of interpretability and helps users understand where the model is focusing during classification. The project report describes the use of Grad-CAM as part of the prediction workflow.

Finally, the study aims to evaluate the performance and reliability of the developed system using appropriate classification measures. Accuracy, precision, recall, F1-score, and confusion matrix analysis are considered for evaluating model performance. The system is intended primarily for preliminary screening and decision support rather than as a substitute for professional dermatological diagnosis. This limitation is consistent with the project's stated purpose of supporting early assessment while retaining the role of clinical expertise.

4.METHODOLOGY

The methodology adopted for Skin DocNet follows a sequential process in which skin lesion images are collected, prepared, validated, classified, and presented through a web-based application. The complete workflow is designed around an EfficientNet-based deep learning model and includes image preprocessing, model prediction, result interpretation, and performance evaluation. The approach is divided into the following stages.

4.1 Dataset Collection and Preparation

The first stage involves preparing a dataset containing images of skin lesions. For the proposed binary classification system, the images are organized into two categories: melanoma and benign lesions. The images are arranged according to their respective class labels so that the model can learn the visual differences between the two categories during training.

Before model training, the image collection is prepared to provide consistent inputs. Images are checked for suitability

and organized systematically. Appropriate preprocessing is applied to make the images compatible with the selected deep learning architecture. The project methodology also considers random ordering and renaming of images to reduce the possibility of unintended learning from file names or image ordering.

4.2 Image Preprocessing

Image preprocessing is performed before an image is supplied to the EfficientNet model. The uploaded image is loaded and converted into an appropriate format before being resized to the required input dimensions. Pixel values are normalized so that the model receives standardized numerical data.

Preprocessing helps maintain consistency between training images and images supplied through the web application. In the implemented system, OpenCV is used for image loading, resizing, normalization, and other image-processing operations.

4.3 Image Validation

An image validation stage is included before classification. When a user uploads an image, the application first checks whether the file is an accepted image type. The system then performs validation to determine whether the uploaded input is suitable for skin lesion analysis. Images that do not satisfy the required conditions are rejected instead of being passed directly to the classification model.

This stage is important because a deep learning classifier may produce an output even when an inappropriate image is supplied. Input validation therefore acts as a protective layer between the user and the prediction model. The project specifically includes validation of uploaded images before prediction.

4.4 EfficientNet Model Development

After preprocessing, the image is provided to the EfficientNet-based classification model. EfficientNet is used as the principal deep learning architecture because it provides a structured balance between model complexity and computational efficiency. The architecture learns useful visual representations from the input images through its deep neural network layers.

During training, the model learns characteristics that help distinguish melanoma from benign lesions. These characteristics can include variations in colour, texture, shape, and other visual patterns present in the lesion images. Training and validation data are used to observe model behaviour and identify potential overfitting. The trained model is subsequently saved for integration with the web application. The project report identifies EfficientNet as the core architecture used for classification.

4.5 Model Prediction

Once a valid image has been preprocessed, it is passed to the trained EfficientNet model for inference. The model generates probability values corresponding to the two target categories.

The category with the higher predicted probability is selected as the final classification.

The prediction stage produces two main outputs: the predicted class and its associated confidence value. These results are then transferred to the application layer so that they can be displayed to the user. The implemented workflow uses a prediction function to obtain the classification, confidence, and probability values from the uploaded image.

4.6 Grad-CAM-Based Interpretation

To make the prediction more understandable, the system incorporates Grad-CAM as an explainability technique. Grad-CAM generates a heatmap that indicates the image regions contributing to the model's classification decision. This allows the prediction to be examined visually rather than relying only on the numerical output.

The generated heatmap is associated with the uploaded image and presented as part of the prediction result. The project implementation includes a separate Grad-CAM generation module after the model prediction stage.

4.7 Web Application Development

The trained model is integrated into a web application developed using Flask. The application provides the interface through which users can register, log in, upload skin lesion images, and view classification results. The backend manages authentication, image processing, model inference, and communication between the different system components.

Flask is used because it provides the application framework required to connect the trained model with the user interface. HTML, CSS, and JavaScript are used to construct the frontend and provide an interactive environment for users.

4.8 Result Visualization

After classification, the prediction information is presented through the result interface. The output can include the predicted category, confidence value, probability information, and Grad-CAM visualization. Presenting these elements together provides both the classification outcome and an indication of the model's visual reasoning.

The system also supports graphical presentation of model-related information. Chart.js is used in the project for displaying probability distributions and performance-related visualizations.

4.9 Performance Evaluation

The final stage involves evaluating the performance of the classification system. Standard classification measures such as accuracy, precision, recall, and F1-score are considered for assessing the model. A confusion matrix is also used to examine correct and incorrect predictions for the melanoma and benign categories.

Performance evaluation helps determine whether the model is producing consistent classification results and identifies areas where misclassification may occur. The project also uses graphical representations and confusion-matrix analysis as part of its evaluation process.

5.LITERATURE SURVEY

The use of deep learning for skin lesion classification has received considerable attention because of the ability of neural networks to learn complex visual patterns from medical images. Researchers have investigated different convolutional neural network architectures, transfer learning strategies, image preprocessing techniques, and hybrid models to improve the identification of skin diseases. The following studies provide the background for the development of the proposed Skin DocNet system.

Alruwaili and Mohamed (2025) investigated an integrated deep learning approach for multi-class skin disease classification by combining EfficientNet-B0, EfficientNet-B2, and ResNet50. Their approach used image augmentation techniques such as rotation, flipping, and brightness and contrast adjustments. The reported accuracy was 99.14%, with precision, recall, and F1-score reaching 0.9914. The study identified external validation, computational efficiency, and broader disease coverage as areas requiring further investigation.

Gulzar et al. (2025) proposed a hybrid transfer learning approach that combined DenseNet121 and EfficientNet-B0 for skin-condition prediction. The approach applied image resizing, normalization, and augmentation operations to improve the training data. The study reported high training and validation performance and highlighted the need for broader disease coverage, mobile deployment, explainable artificial intelligence, and validation across different clinical environments. The use of EfficientNet in this work demonstrates the suitability of efficient transfer-learning architectures for medical image classification.

Aboulmira et al. (2025) developed a hybrid model that combined EfficientNet-B0 with wavelet decomposition for skin cancer classification. The method incorporated image enhancement, colour normalization, rotation, flipping, and multi-level wavelet processing. The study reported strong classification performance on skin-lesion datasets and suggested reducing model complexity while improving performance on minority classes as future research directions.

Wei et al. (2023) presented a skin disease classification model based on the fusion of DenseNet and ConvNeXt architectures. Their approach incorporated normalization, image resizing, and advanced augmentation methods such as Cutout, CutMix, and AugMix. The study demonstrated that combining different feature-learning architectures can improve classification performance. However, lightweight optimization, broader disease coverage, and cross-dataset validation remained important areas for future development.

Sadik et al. (2023) conducted an extensive investigation of convolutional neural network architectures using transfer learning for skin-condition diagnosis. The study considered architectures including MobileNet, Xception, ResNet50, and

InceptionV3 together with image augmentation techniques. Transfer learning combined with augmentation produced strong classification results, demonstrating the usefulness of pretrained CNN architectures for skin-image analysis.

Hoang et al. (2022) investigated lightweight deep learning for multi-class skin-lesion classification using Wide-ShuffleNet and EfficientNet-B0. Their methodology included segmentation, filtering, and binarization before classification. The reported performance on HAM10000 and ISIC datasets demonstrated the potential of lightweight networks for skin-lesion analysis. The study also identified mobile deployment and class-imbalance handling as important areas for further research.

Muhaba et al. (2022) studied automatic skin-disease diagnosis using clinical images and patient information. Their approach incorporated colour-constancy processing and image augmentation and employed MobileNet-V2 for classification. The study demonstrated that combining image-based learning with clinical information can support automated skin-disease assessment and highlighted real-world clinical decision support as an important application area.

Almuayqil, Abd El-Ghany, and Elmogy (2022) proposed a hybrid deep learning framework for multi-type skin-disease diagnosis. Their approach used multiple CNN feature extractors, including VGG19, InceptionV3, ResNet50, DenseNet201, Xception, and InceptionResNet. The reported results showed high classification performance, while the authors identified dimensionality reduction, larger datasets, and real-time multi-disease systems as potential directions for further development.

Kshirsagar et al. (2022) investigated deep learning methods for automated skin-disease prognosis using image enhancement, segmentation, normalization, and multi-scale processing. Their hybrid approach combined MobileNetV2 and LSTM with fusion-based segmentation. The study demonstrated that preprocessing and feature-fusion strategies can contribute to skin-disease classification, while improvements in generalization and handling of different lighting conditions remain important.

Anand et al. (2022) examined automated skin-disease diagnosis using dermoscopic images. Their work proposed a CNN-based approach and compared its performance with ResNet18, ResNet50, and ResNet101. Image normalization, augmentation, cropping, and contrast correction were used during preprocessing. The study demonstrated the effectiveness of CNN models for automated lesion classification and suggested real-time deployment and explainable approaches as future research opportunities.

Alam et al. (2021) developed SkinDiseaseNet, a lightweight CNN architecture for skin-disease identification. The study used image resizing, normalization, and augmentation and reported an accuracy of 92.87%. The work demonstrated that lightweight CNN models can provide useful classification performance while reducing computational requirements, making them suitable candidates for practical deployment.

Zhu, Wang, and Chen (2021) proposed a framework for diagnosing different skin conditions using a modified EfficientNet-B4 architecture with auxiliary classifiers. The images were resized and augmented using rotation and flipping, together with colour-constancy normalization. The study reported an accuracy of 94.8%, sensitivity of 93.4%, specificity of 95.0%, and AUC of 0.985. The authors identified dataset diversity and multi-centre validation as areas requiring additional investigation.

Bajwa et al. (2020) evaluated several deep neural network architectures for computer-aided skin-disease diagnosis. The investigated architectures included ResNet-152, DenseNet-161, SE-ResNeXt-101, and NASNet. The study reported different performance levels across DermNet and ISIC datasets and emphasized the importance of larger datasets, real-time computer-aided diagnosis, and demographic generalization.

Wu et al. (2020) investigated automated diagnosis of inflammatory skin conditions using deep learning. Their approach used image resizing, rotation, and flipping with a modified EfficientNet-B4 architecture. The reported overall accuracy was 95.80%, demonstrating the potential of EfficientNet-based approaches for dermatological image classification. Smartphone-based diagnosis and integration with clinical testing were identified as possible future directions.

6. SYSTEM DESIGN




Figure 1: System Architecture

Figure 1 illustrates the overall system architecture of the proposed Skin DocNet application. The architecture shows how a skin lesion image moves through different components of the system, starting from user interaction and ending with the classification result.

The process begins with the User, who registers or logs into the application and uploads a skin lesion image through the Web Interface. The Web Interface transfers the user request to the Backend Application, which controls the main operations of the system. The backend also communicates with the User Database, which maintains user credentials and session information.

After receiving the uploaded image, the Backend Application forwards it to the Image Preprocessing and Validation module. At this stage, the image is checked and prepared for analysis. The processed image is then passed to the CNN Model, which performs image classification and disease prediction. In the proposed Skin DocNet system, the CNN-based approach uses EfficientNet for skin lesion classification.

Once the model completes the prediction, the output is sent to the Result Display module. The result contains the predicted disease category and confidence information, along with other supporting information provided by the application. Thus, the figure represents the complete flow of Skin DocNet from user input to image processing, model-based prediction, and final result presentation. The same architecture in the project report describes the web interface, backend, database, preprocessing, CNN model, and result display as the main interacting components.

7.SCREENSHOTS

1.Home Page



Figure 2: Home Page

2.Login Page



Figure 3: Login Page

3.Image Upload Page




Figure 4: Upload Page

4.Prediction Result Page



Figure 5: Prediction Result Page

5.Model accuracy page



Figure 6: Model accuracy page

8.CONCLUSION AND FUTURE SCOPE

The proposed Skin DocNet system demonstrates the application of deep learning for automated melanoma and benign skin lesion classification using the EfficientNet architecture. The developed approach integrates image preprocessing, input validation, deep learning-based classification, and web-based result visualization into a unified framework. The system processes an uploaded skin lesion image and generates a classification result with associated confidence information. Grad-CAM is incorporated to provide a visual interpretation of the model's decision, thereby improving the transparency of the prediction process. The overall framework demonstrates the potential of combining efficient deep learning architectures with accessible web technologies for preliminary skin lesion assessment.

Although the proposed system provides a useful foundation for computer-assisted skin lesion screening, further research is required before it can be considered for clinical use. Future work can focus on expanding the dataset with more diverse images and increasing the number of lesion categories. Improving model generalization across different imaging conditions and patient populations is also important. The system can additionally be extended to mobile platforms for image-based assessment using smartphone cameras. Advanced explainable-AI methods and ensemble learning techniques may be investigated to improve prediction reliability and interpretability. Integration with hospital information systems and cloud-based deployment could further improve accessibility, scalability, and information exchange. Real-time image or video analysis can also be explored as an extension of the current static-image classification framework.

REFERENCES

- [1] M. Alruwaili and M. Mohamed, "An Integrated Deep Learning Model with EfficientNet and ResNet for Accurate Multi-Class Skin Disease Classification," 2025.
- [2] Y. Gulzar, S. Agarwal, S. Soomro, M. Kandpal, C. W. Onn, S. Saini, and A. Bounsiar, "Next-generation Approach to Skin Disorder Prediction Employing Hybrid Deep Transfer Learning," 2025.
- [3] A. Aboulmira, H. Hrimech, M. Lachgar, M. Hanine, C. O. Garcia, G. Mezquita, and I. Ashraf, "Hybrid Model with Wavelet Decomposition and EfficientNet for Accurate Skin Cancer Classification," 2025.
- [4] M. Wei, Q. Wu, H. Ji, J. Wang, T. Lyu, J. Liu, and L. Zhao, "A Skin Disease Classification Model Based on DenseNet and ConvNeXt Fusion," 2023.
- [5] R. Sadik, A. Majumder, A. A. Biswas, B. Ahammad, and M. M. Rahman, "An In-Depth Analysis of Convolutional Neural Network Architectures With Transfer Learning for Skin Disease Diagnosis," 2023.
- [6] L. Hoang et al., "Classification of Multiclass Skin Lesions using Lightweight Deep Learning," 2022.
- [7] K. A. Muhaba et al., "Deep Learning-Based Automatic Skin Disease Diagnosis from Clinical Images and Patient Data," 2022.
- [8] S. N. Almuayqil, S. Abd El-Ghany, and M. Elmogy, "A Hybrid Deep Learning Model for Feature Fusion in Multi-Type Computer-Aided Diagnosis for Early Skin Disease Signs," 2022.
- [9] P. R. Kshirsagar, H. Manoharan, S. Shitharth, and A. M. Abdulrhman, "Deep Learning Methods for Automated Skin Disease Prognosis," 2022.
- [10] V. Anand, S. Gupta, D. Koundal, and S. Mahajan, "Automated Skin Disease Diagnosis Using Dermoscopy and Deep Learning," 2022.
- [11] M. A. Alam et al., "Lightweight CNN Model for Skin Disease Identification: SkinDiseaseNet," 2021.
- [12] C.-Y. Zhu, Y.-K. Wang, and H.-P. Chen, "A Framework for Diagnosing Various Skin Conditions in a Clinical Setting Using Deep Learning," 2021.
- [13] M. N. Bajwa et al., "Deep Neural Networks for Computer-Aided Skin Disease Diagnosis," 2020.
- [14] H. Wu et al., "An Automated Diagnosis for Inflammatory Skin Conditions Using Deep Learning Images," 2020.