



International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 15, Issue 3, March 2026

Automated Skin Disease Detection Using Transfer Learning with Mobile Net CNN

Vaishali Sanap, Harshada Barve, Sharvari Deshmukh, Gitanjalee Funde, Jiya Jori

Lecturer, Dept. of Computer Engineering, AISSMS Polytechnic, Pune, India

Student, Dept. of Computer Engineering, AISSMS Polytechnic, Pune, India

Student, Dept. of Computer Engineering, AISSMS Polytechnic, Pune, India

Student, Dept. of Computer Engineering, AISSMS Polytechnic, Pune, India

Student, Dept. of Computer Engineering, AISSMS Polytechnic, Pune, India

ABSTRACT: Skin diseases are among the most prevalent health conditions worldwide, yet early and accurate diagnosis remains challenging due to limited access to dermatological expertise, particularly in rural and underdeveloped regions [1][3]. Traditional diagnosis methods are time-consuming, subjective, and dependent on specialist experience. This research proposes an automated skin disease detection system using Deep Learning, specifically Convolutional Neural Networks (CNNs) with a MobileNet architecture, to classify skin disease images efficiently and accurately. Transfer learning is employed to leverage pre-trained MobileNet weights, enabling improved performance with reduced computational cost[8]. The system includes image preprocessing, augmentation, model training, and confidence-based prediction output[5]. Experimental results demonstrate that the MobileNet-based approach achieves high classification accuracy, faster inference time, and reduced overfitting compared to a baseline CNN model.

KEYWORDS: Deep Learning, Convolutional Neural Networks, MobileNet, Transfer Learning, Skin Disease Classification, Medical Image Processing

I. INTRODUCTION

Skin diseases represent one of the most widespread categories of health disorders worldwide, affecting individuals of all ages and demographics [1][3]. Conditions such as acne,

eczema, psoriasis, vitiligo, fungal infections, and other dermatological abnormalities not only cause physical discomfort but can also lead to psychological stress, reduced self-confidence, and long-term medical complications if left untreated[1]. Early and accurate diagnosis is therefore essential to prevent disease progression and ensure timely treatment.

Despite the importance of early detection, access to qualified dermatologists remains limited in many rural and underdeveloped regions [9]. Patients often rely on general practitioners or self-diagnosis, which may result in delayed treatment or misdiagnosis. Furthermore, manual clinical

examination is inherently subjective and depends heavily on the experience and expertise of the medical professional. Variations in lighting conditions, skin tone, lesion appearance, and human judgment can influence diagnostic consistency.

With the rapid advancement of Artificial Intelligence (AI) and Machine Learning (ML), automated diagnostic systems have emerged as promising tools in healthcare. Among these, Deep Learning techniques—particularly Convolutional Neural Networks (CNNs)—have demonstrated exceptional performance in medical image analysis[5][6]. CNNs are specifically designed to process image data by automatically learning hierarchical features such as edges, textures, shapes, and complex visual patterns without requiring manual feature engineering. This capability makes them highly suitable for dermatological image classification tasks [5].

However, traditional deep CNN architectures often require high computational power and large memory resources, making them challenging to deploy in real-time or resource-constrained environments[15]. MobileNet is a lightweight model that reduces computation while maintaining good accuracy, making it suitable for mobile applications.

This project proposes an automated skin disease detection system using a MobileNet-based CNN with transfer learning to classify diseases from images and provide predictions with confidence scores.

II. RELATED WORK

I. **Y. Qian et al. [1]** analyzed racial disparities in melanoma using the SEER database (1975–2016). Their study highlighted significant variations in disease outcomes across different racial groups, emphasizing the need for more inclusive and accessible diagnostic systems.

II. Similarly, **S. Babool et al. [2]** conducted a systematic review focusing on skin tone representation in dermatomyositis rashes. Their findings revealed a lack of diversity in dermatological datasets, which can negatively impact the performance of AI- based diagnostic models.

III. **A. F. Alexis et al. [3]** performed a comparative survey on common dermatological disorders in individuals with different skin tones. The study emphasized the importance of accurate diagnosis across diverse populations and highlighted challenges faced in clinical practice.

IV. In the field of deep learning, **M. Frid- Adar et al. [4]** proposed the use of Generative Adversarial Networks (GANs) for synthetic medical image augmentation. Their approach improved CNN performance by increasing the diversity of training data, particularly in cases with limited datasets.

V. **C. Shorten and T. M. Khoshgoftaar [5]** presented a comprehensive survey on image data augmentation techniques used in deep learning. The study discussed various methods to improve model generalization and reduce overfitting, especially in small datasets.

VI. **A. Radford et al. [6]** Furthermore, introduced Deep Convolutional Generative Adversarial Networks (DCGANs), which demonstrated the capability of unsupervised learning in generating high-quality images. This work laid the foundation for applying GANs in medical imaging tasks.

III. DATABASE

The dataset used in this study was obtained from the publicly available Kaggle repository, which contains labeled images of various dermatological conditions. The dataset includes images representing six different skin categories: Acne, Eczema, Psoriasis, Vitiligo, Tinea, and Normal Skin. Each image in the dataset is associated with a predefined class label, enabling supervised learning for the skin disease classification task.

The dataset was organized in a folder- based directory structure, where each folder corresponds to a specific skin disease category. This structure allows deep learning frameworks to automatically assign class labels during the training process. Since the dataset was already curated and labeled by the dataset contributor, manual annotation was not required.

Before training the model, the dataset was carefully inspected to ensure image integrity. Any corrupted or unreadable image files were removed to maintain dataset quality. The dataset was then divided into training and testing subsets using an 80:20 split. The training subset was used to train the deep learning model, while the testing subset was used to evaluate the model's performance on previously unseen data.

To ensure compatibility with the MobileNet architecture, all images were resized to 224×224 pixels, which corresponds to the standard input size required by the model. Additionally, pixel values were normalized during preprocessing to improve training stability and model convergence[5].

Although the dataset was partially preprocessed, additional preprocessing steps such as resizing and normalization were applied during the training pipeline to ensure consistent input representation. These preprocessing steps help improve the robustness and accuracy of the classification model. The distribution of images across the six skin disease categories used in this research is shown in Table I.

DATASET CLASS DISTRIBUTION

Class	Number of Images
Acne	600
Eczema	500
Psoriasis	480
Vitiligo	450
Tinea	420
Normal Skin	550

Table I illustrates the number of images available for each skin disease category. The dataset contains a diverse set of dermatological images with variations in lighting conditions, skin tones, and lesion characteristics. This diversity enables the model to learn meaningful visual patterns associated with different skin diseases and improves its ability to generalize when classifying unseen images.

IV. PROPOSED METHODOLOGY

4.2 Test Images

Figure 4.1 shows sample test images from different skin disease categories used to evaluate the trained model. These images represent real-world variations in skin texture, color, lighting conditions, and lesion patterns. The test samples include conditions such as acne, eczema, psoriasis, vitiligo, tinea, and normal skin. The purpose of presenting these test images is to demonstrate the diversity of the dataset and to validate that the trained MobileNet model can generalize well to unseen images. Despite variations in background and skin tone, the model successfully extracted discriminative features and produced accurate.



4.2 Accuracy comparison between CNN and MobileNet

Baseline CNN Model

Initially, a Convolutional Neural Network model was implemented as a baseline model for skin disease classification. CNNs are widely used in image classification tasks because they can automatically extract spatial features from images using convolutional filters.

The CNN architecture used in this study consists of multiple layers including:

- Convolution layers

- Activation functions
- Pooling layers
- Fully connected layers

These layers work together to detect features such as edges, textures, and patterns in the images. However, during experimentation, the CNN model achieved relatively lower accuracy compared to the expected performance for medical image **classification tasks**. **This limitation** motivated the use of a more efficient deep learning architecture.

MobileNet Architecture

To improve the performance of the skin disease classification system, the MobileNet architecture was used. MobileNet is a lightweight convolutional neural network designed for efficient image classification. It uses depthwise separable convolutions, which split the operation into depthwise and pointwise convolutions, reducing computational cost and the number of parameters.

In this work, transfer learning was applied by adapting a pre-trained MobileNet model for six skin disease classes. The final classification layer was modified accordingly. The Softmax activation function is used in the output layer to generate probability values, and the class with the highest probability is selected as the final prediction.

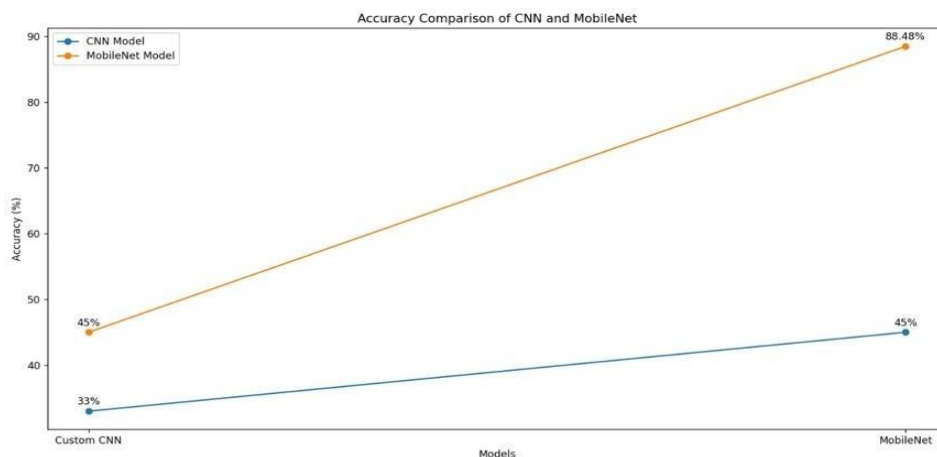


Fig 4.2: Accuracy Comparison of CNN and MobileNet

4.3 System Architecture and Working principle

The overall architecture of the proposed skin disease detection system is illustrated in **Figure 4.1**. The system is designed to automatically analyze dermatological images and identify different skin conditions using a deep learning approach. The architecture integrates multiple processing stages that work together to transform raw skin images into meaningful diagnostic predictions. By combining image processing techniques with a convolutional neural network, the system aims to improve the efficiency and reliability of automated skin disease classification.

The proposed framework uses a deep learning model based on the **MobileNetV2 architecture**, which is known for its lightweight design and strong performance in image classification tasks. It efficiently extracts important visual features from images while requiring low computational resources, making it suitable for accurate medical image analysis.

The architecture follows a sequential process where the input skin image is first preprocessed and augmented, then passed to a deep learning model for feature extraction and classification. Each stage helps standardize the data and capture important visual features. The final output is the predicted skin disease along with a confidence score indicating the probability of the prediction.

1. Input Skin Image

The process begins with the acquisition of a skin image. The user uploads an image of the affected skin area through the system interface. The input image acts as the primary data source for analysis. These images may be captured using a smartphone camera or other imaging devices. The quality and clarity of the input image play an important role in improving the accuracy of the prediction.

2. Image Preprocessing

After the image is uploaded, it undergoes preprocessing to prepare it for further analysis. Image preprocessing involves operations such as resizing and normalization. Resizing ensures that all input images have the same dimensions required by the deep learning model. Normalization adjusts pixel values to a standard range, which helps improve model stability and training performance.

3. Data Augmentation

To improve the generalization capability of the model and increase dataset diversity, data augmentation techniques are applied. Data augmentation artificially expands the dataset by generating modified versions of existing images. In this system, augmentation operations such as rotation, flipping, zooming, brightness adjustment, and noise addition are applied. These transformations help the model learn robust features and reduce overfitting during training.

4. MobileNetV2 Feature Extraction

The augmented images are then passed to the MobileNetV2 architecture for feature extraction. MobileNetV2 is a lightweight convolutional neural network designed for efficient image classification tasks. It utilizes depthwise separable convolutions and inverted residual blocks to extract meaningful features from images while maintaining computational efficiency. In this stage, the model identifies important patterns such as texture, color variation, and lesion structure present in the skin images.

5. Fully Connected Layer

After feature extraction, the extracted feature maps are passed to a fully connected layer. This layer combines the extracted features and prepares them for the final classification stage. The fully connected layer acts as a classifier that learns the relationship between extracted features and their corresponding disease categories.

6. Softmax Classification

The final classification is performed using a Softmax activation function. The Softmax layer converts the output values into probability scores for each skin disease category. This allows the model to determine the likelihood of the input image belonging to each class.

7. Predicted Skin Disease Output

In the final stage, the system outputs the predicted skin disease based on the highest probability score generated by the Softmax classifier. The system categorizes the skin condition into classes such as melanoma, eczema, psoriasis, acne, or other skin conditions. The prediction result is displayed to the user, helping them understand the possible skin condition detected by the model.

8. MobileNetV2 Feature Extraction

The augmented images are then passed to the MobileNetV2 architecture for feature extraction. MobileNetV2 is a lightweight convolutional neural network designed for efficient image classification tasks. It utilizes depthwise separable convolutions and inverted residual blocks to extract meaningful features from images while maintaining computational efficiency. In this stage, the model identifies important patterns such as texture, color variation, and lesion structure present in the skin images.

9. Fully Connected Layer

After feature extraction, the extracted feature maps are passed to a fully connected layer. This layer combines the extracted features and prepares them for the final classification stage. The fully connected layer acts as a classifier that learns the relationship between extracted features and their corresponding disease categories.

10. Softmax Classification

The final classification is performed using a Softmax activation function. The Softmax layer converts the output values into probability scores for each skin disease category. This allows the model to determine the likelihood of the input image belonging to each class.

11. Predicted Skin Disease Output

In the final stage, the system outputs the predicted skin disease based on the highest probability score generated by the Softmax classifier. The system categorizes the skin condition into classes such as melanoma, eczema, psoriasis, acne, or other skin conditions. The prediction result is displayed to the user, helping them understand the possible skin condition detected by the model.

4.4 Model Training

The skin disease classification model was developed using transfer learning with the MobileNetV2 deep learning architecture implemented in the TensorFlow and Keras frameworks. Transfer learning was used to leverage feature representations learned from large-scale datasets, allowing efficient training with a relatively smaller dermatological image dataset.

The MobileNetV2 model was initialized with pre-trained weights from the ImageNet dataset while excluding the original classification layers. This allows the model to retain general image feature extraction capabilities while adapting to the specific skin disease classification task. To perform classification for the six skin disease categories, additional layers were added on top of the base MobileNetV2 model. These layers include a Global Average Pooling layer, followed by a fully connected dense layer with 256 neurons and ReLU activation, and a Dropout layer with a dropout rate of 0.5 to reduce overfitting. Finally, a Softmax output layer with six neurons was used to generate probability scores for each skin disease class.

During the initial training phase, the base MobileNetV2 layers were frozen so that only the newly added classification layers were trained. This allows the model to adapt the final layers to the skin disease dataset without modifying the pre-trained feature extraction layers. The model was trained using the Adam optimizer with a learning rate of 0.0001, and the categorical cross-entropy loss function was used for multi-class classification. The training process used a batch size of 32 images and was performed for 30 epochs. To improve model generalization and reduce overfitting, several data augmentation techniques were applied using the ImageDataGenerator utility in Keras. These techniques included rotation, zooming, horizontal flipping, width and height shifting, and shear transformations.

After the initial training stage, fine-tuning was performed to further improve model performance. In this stage, the last 20 layers of the MobileNetV2 base model were unfrozen, allowing their weights to be updated during training. Fine-tuning was conducted with a smaller learning rate of 0.00001 to ensure stable optimization without disrupting previously learned features. The model was then trained for an additional 10 epochs. Finally, the trained model was evaluated using the testing dataset to measure its classification performance. The evaluation results demonstrated that the MobileNetV2-based model achieved high classification accuracy for the skin disease detection task while maintaining efficient computational performance.

V. RESULT

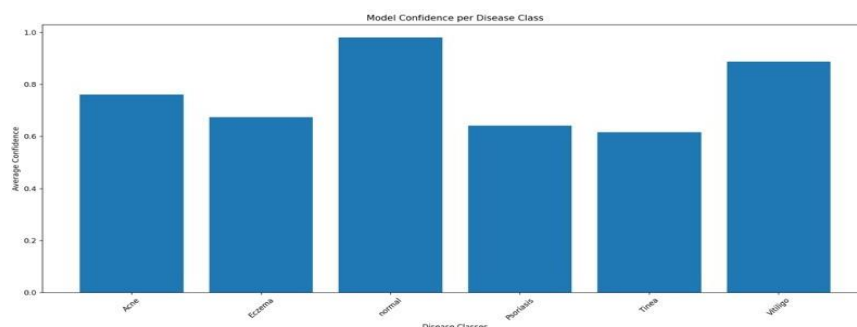


Fig 5.1: Average Confidence per Skin Disease Class

To evaluate the performance of the proposed skin disease detection system, multiple test images from different categories were provided as input to the trained model. The system successfully generated predictions along with corresponding confidence scores.

From the graph, it can be observed that the model achieves varying confidence levels across different classes. The highest confidence is observed for **Normal skin (~0.98)**, indicating that the model can very accurately distinguish healthy skin from diseased conditions. This is followed by **Vitiligo (~0.88)** and **Acne (~0.76)**, showing strong prediction capability for these classes.

Moderate confidence scores are seen for **Eczema (~0.67)**, **Psoriasis (~0.64)**, and **Tinea (~0.61)**.

These slightly lower values may be due to visual similarities between these conditions, making classification more challenging. Overall, the model demonstrates consistent and reliable performance across all classes. The confidence scores represent the probability of the predicted class, reflecting the model's certainty in its decisions.

The results indicate that the use of the MobileNetV2 architecture provides a good balance between accuracy and computational efficiency. The model generalizes well across different skin conditions and is capable of handling real-world variations in images. Additionally, the visual output of the system, including predicted labels and confidence scores, enhances usability by providing clear and interpretable results. This makes the system effective not only for automated detection but also as a supportive tool for understanding potential skin conditions.

VI. CONCLUSION

This research presented a deep learning- based system for automated skin disease classification using dermatological images. The system was developed using a dataset obtained from Kaggle containing six different skin disease categories.

Initially, a CNN model was implemented as a baseline approach, but it showed relatively low accuracy. To improve classification performance, the MobileNet architecture was implemented due to its efficient design and strong feature extraction capability. Experimental results demonstrate that the MobileNet model significantly improves classification accuracy compared to the CNN model. The system can successfully predict skin diseases and provide confidence scores for each prediction. The proposed approach highlights the potential of deep learning techniques in assisting dermatological diagnosis and improving accessibility to healthcare services.

Future work may include expanding the dataset, improving model accuracy using advanced deep learning techniques, and integrating the system into mobile or web- based healthcare platforms for real-time diagnosis.

REFERENCES

1. **Y. Qian, P. Johannet, A. Sawyers, J. Yu, I. Osman, and J. Zhong**, "The ongoing racial disparities in melanoma: An analysis of the surveillance, epidemiology, and end results database (1975-2016)," *Journal of the American Academy of Dermatology*, vol. 84, no. 6, pp. 1585–1593, 2021.
2. **S. Babool, S. F. Bhai, C. Sanderson, A. Salter, and L. Christopher Stine**, "Racial disparities in skin tone representation of dermatomyositis rashes: a systematic review," *Rheumatology*, vol. 61, no. 6, pp. 2255–2261, 2022.
3. **A.F.Alexis, A. B. Sergay, and S. C. Taylor**, "Common dermatologic disorders in skin of color: a comparative practice survey," *Cutis*, vol. 80, no. 5, pp. 387–394, 2007.
4. **M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan**, "Gan-based synthetic medical image augmentation for increased cnn performance in liver lesion classification," *Neurocomputing*, vol. 321, pp. 321–331, 2018.
5. **C. Shorten and T. M. Khoshgoftaar**, "A survey on image data augmentation for deep learning," *Journal of big data*, vol. 6, no. 1, pp. 1–48, 2019.
6. **A. Radford, L. Metz, and S. Chintala**, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
7. **M. Groh, C. Harris, L. Soenksen, F. Lau, R. Han, A. Kim, A. Koochek, and O. Badri**, "Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 1820–1828.

8. **M. Groh, C. Harris, R. Daneshjou, O. Badri, and A. Koochek**, "Towards transparency in dermatology image datasets with skin tone annotations by experts, crowds, and an algorithm," arXiv preprint arXiv:2207.02942, 2022.
9. **G. Kleinberg, M. J. Diaz, S. Batchu, and B. Lucke-Wold**, "Racial underrepresentation in dermatological datasets leads to biased machine learning models and inequitable healthcare," *Journal of biomed research*, vol. 3, no. 1, p. 42, 2022.
10. **L. Gao, L. Zhang, C. Liu, and S. Wu**, "Handling imbalanced medical image data: A deep-learning-based one-class classification approach," *Artificial Intelligence in Medicine*, vol. 108, p. 101935, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S09333657203026>
11. **J. M. Johnson and T. M. Khoshgoftaar**, "Survey on deep learning with class imbalance," *Journal of Big Data*, vol. 6, no. 1, pp. 1–54, 2019.
12. **X. Yi, E. Walia, and P. Babyn**, "Generative adversarial network in medical imaging: A review," *Medical Image Analysis*, vol. 58, p. 101552, Dec. 2019. [Online]. Available: <http://dx.doi.org/10.1016/j.media.2019.10.1552>
13. **Y. Hong, U. Hwang, J. Yoo, and S. Yoon**, "How generative adversarial networks and their variants work: An overview," *ACM Computing Surveys (CSUR)*, vol. 52, no. 1, pp. 1–43, 2019.
14. **P. Tschandl, C. Rosendahl, and H. Kittler**, "The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions," *Scientific data*, vol. 5, no. 1, pp. 1–9, 2018.
15. **K. He, X. Zhang, S. Ren, and J. Sun**, "Deep residual learning for image recognition," 2015.
16. **I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde Farley, S. Ozair, A. Courville, and Y. Bengio**, "Generative adversarial networks," 2014.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details