International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 06 | June - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

# Automatic Detection of Skin Lesions Using CNN And Mobilenetv2

1. Dr. P. Rajendra Prasad
Associate Professor, Department of Computer Science and Engineering Vignan's Institute of Management and Technology for Women, Hyd.
Email: rajipe@gmail.com

2. .Pemira Chandrakala
UG Student, Department of Computer Science and Engineering Vignan's Institute of Management and Technology for Women, Hyd.
Email: pemirachandrakala042@gmail.com

3. Dendi Umabindu
UG Student, Department of Computer Science and Engineering Vignan's Institute of Management and Technology for Women, Hyd.
Email: dendiumabindu@gmail.com

4. Erra Meghana
UG Student, Department of Computer Science and Engineering Vignan's Institute of Management and Technology for Women, Hyd.
Email: meghanaerra76@gmail.com

**Abstract—Skin diseases are among the most common health problems worldwide, often requiring early and accurate diagnosis to prevent severe outcomes. This study presents a deep learning-based dermatological disease detection system using MobileNetV2, a lightweight convolutional neural network. We utilize the HAM10000 dataset containing dermatoscopic images of seven different types of skin lesions. Our method incorporates advanced preprocessing techniques, hair artifact removal (Dull Razor), and data augmentation to improve the model's performance. MobileNetV2 is finetuned through transfer learning and trained on pre processed data to classify skin conditions efficiently. The proposed model achieved significant classification accuracy while maintaining computational efficiency, making it suitable for deployment in real-world clinical and mobile applications.**

**Keywords - Deep learning, MobileNetV2, dermatological disease detection, skin cancer classification, HAM10000 dataset, convolutional neural networks (CNN), image preprocessing, transfer learning**

## I. INTRODUCTION

Skin diseases are among the most common health concerns affecting people worldwide, ranging from mild irritations to life-threatening conditions such as melanoma. The increasing incidence of skin-related disorders, particularly skin cancers, highlights the urgent need for reliable and early diagnostic methods. Melanoma, for instance, is known for its aggressive behaviour and high mortality rate if not identified and treated at an early stage. However, these methods can be subjective, time-consuming, and prone to interobserver variability, especially when the differences between benign and malignant lesions are subtle with the rapid advancement of artificial intelligence and machine learning technologies, deep learning has emerged as a powerful tool for image-based medical diagnosis. Specifically, Convolutional Neural Networks (CNNs) have demonstrated exceptional capabilities in image classification and object detection tasks, making them suitable for analysing dermoscopic images. In recent years, several CNN based models have been developed for skin lesion classification, often achieving dermatologist-level performance. Despite their accuracy, many of these models require significant computational resources, which limits their application in real-time or resource constrained environments.

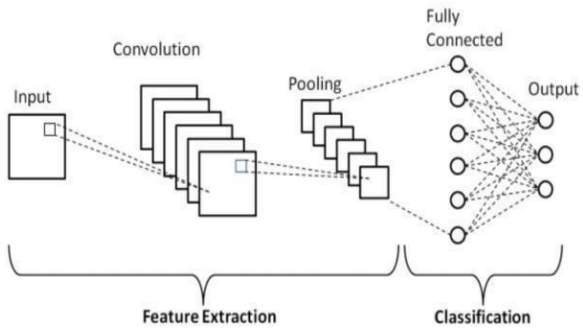


Fig: General structure of a Convolutional Neural Network

To address this limitation, our project focuses on leveraging MobileNetV2, a lightweight and efficient deep learning architecture designed for mobile and embedded vision applications. MobileNetV2 offers an optimal balance between accuracy and computational efficiency, making it an ideal choice for real-time dermatological analysis, especially in telemedicine and

rural healthcare settings. This research involves training the MobileNetV2 model on the HAM10000 dataset, a widely recognized collection of dermoscopic images representing various types of skin lesions. To improve model performance, we employed several preprocessing techniques, such as hair artifact removal, image normalization, and data augmentation. By integrating these methods with transfer learning and fine-tuning strategies, we developed a robust and efficient model capable of detecting multiple skin conditions from dermoscopic images. The primary objective of this work is to develop a deep learning-based system that can support dermatologists in diagnosing skin diseases more accurately and efficiently. The proposed approach aims not only to reduce diagnostic errors but also to enable early detection and treatment of serious skin conditions, thereby improving patient outcomes and accessibility to dermatological care.

## II. LITERATURE REVIEW

Adoption of deep learning for skin disease diagnosis has gained widespread recognition due to its accuracy and automation capabilities. One of the critical resources enabling this progress is the HAM10000 dataset developed by Tschandl et al. [1], which contains over 10,000 annotated dermoscopic images from multiple sources. This dataset serves as a reliable foundation for training and benchmarking classification models in dermatological research. Esteva et al. [2] pioneered a deep learning approach capable of detecting skin cancer with an accuracy comparable to that of experienced dermatologists. Their study showcased how convolutional neural networks (CNNs) could outperform traditional diagnostic tools when trained on high-quality data. Further contributing to this domain, Codella et al. [3] organized the ISIC 2018 challenge, encouraging the development of innovative algorithms for lesion segmentation and classification. Their initiative provided a standard platform for performance comparison and promoted community-driven advancements Preprocessing plays a crucial role in improving image quality for CNN-based models. Abbas et al. [4] presented a comparative study on various hair removal techniques for dermoscopy, helping to reduce visual noise and enhance lesion visibility in input images. To make deep learning models accessible for real-time and mobile applications, Howard et al. [5] introduced the Mobile Net architecture, which reduces the number of trainable parameters through depth wise separable convolutions. This lightweight model architecture enables efficient inference in resource-constrained environments without significant loss in accuracy. Brinker et al. [6] evaluated the performance of a CNN trained with dermoscopic images and found it to be on par with the diagnostic accuracy of over 140 dermatologists. Their work highlights the growing reliability of machine learning in assisting or even automating parts of clinical decision-making. Transfer learning has also played an essential role in medical image analysis. Hosny et al. [7] used pre-trained CNNs to classify skin cancer images, significantly improving accuracy and reducing training time by transferring learned knowledge from general to domain-specific tasks. To further improve performance, Gessert et al. [8] proposed the incorporation of attention mechanisms within CNNs. Their model dynamically focused on the most informative parts of skin lesions, enhancing feature representation and leading to more precise classification outcomes.

## III. METHODOLOGY

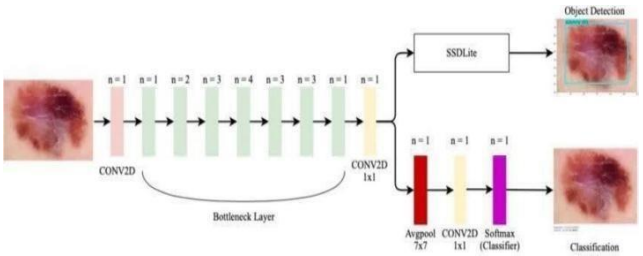


Fig: MobileNetV2 Architecture

The methodology adopted in this research is centred around designing robust, accurate, and computationally efficient system for classifying dermatological diseases from dermoscopic images. The process is structured into key stages, including data acquisition, preprocessing, model design and training, evaluation, and optimization.

### A. Data Acquisition :

It consists of 10,015 dermoscopic images collected from various sources and categorized into seven distinct skin lesion types:

1. Actinic keratoses and intraepithelial carcinoma (akiec)
2. Basal cell carcinoma (bcc)
3. Benign keratosis-like lesions (bkl)
4. Dermatofibroma (df)
5. Melanocytic nevi (nv)
6. Melanoma (mel)
7. Vascular lesions (vasc)

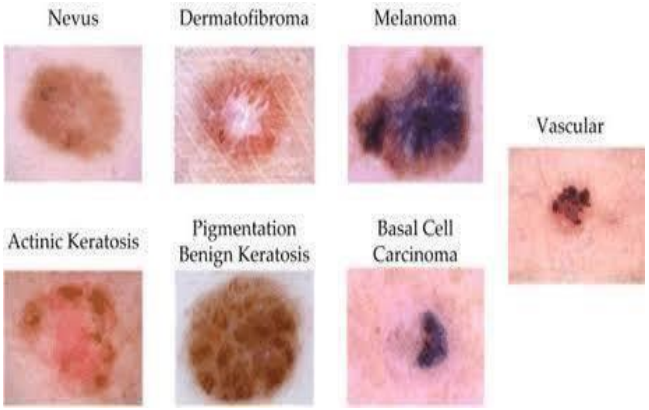

Each image in the dataset is accompanied by metadata and has been clinically confirmed, making it suitable for training deep learning models in a medical context.

**B. Data Preprocessing:**

Preprocessing is an essential step in ensuring that the input data is clean, balanced, and suitable for deep learning models.

1. Image Resizing: All images were resized to a standard resolution of 224x224 pixels to match the input size requirements of the MobileNetV2 model.

2, Normalization: This step helps in faster convergence during training.

3. Class Balancing: The original dataset is imbalanced, with some lesion types significantly underrepresented. To address this, we applied oversampling techniques to replicate minority class images and under sampling for overrepresented classes where appropriate

4. Data Augmentation: To prevent overfitting and improve the generalization capability of the model, various augmentation techniques were employed, such as horizontal and vertical flipping, random rotations, zooming, and brightness adjustments.

5. Artifact Removal: Dermoscopic images often contain artifacts such as hairs and air bubbles. The Dull Razor algorithm was used to detect and remove hair artifacts by identifying thin, dark structures and replacing them using interpolation and filtering techniques.

**C. Model Selection: MobileNetV2**

To build an efficient and lightweight system, we chose MobileNetV2, a deep convolutional neural network optimized for mobile and embedded vision applications. MobileNetV2 introduces inverted residual blocks and depth wise separable convolutions, significantly reducing the number of parameters and computation required compared to traditional CNNs, while maintaining high accuracy.

**D. Transfer Learning**

Since training a deep model from scratch requires a large amount of data and computational resources, we employed transfer learning by using a MobileNetV2 model pre-trained on the ImageNet dataset. The final layers were modified to suit our classification task: The base layers of MobileNetV2 were frozen to retain previously learned low-level features. Additional fully connected layers were added, including a global average pooling layer, a dropout layer to reduce overfitting, and a dense softmax output layer with seven nodes corresponding to the lesion categories.

**E. Training Configuration**

The model was trained using the following settings:
Optimizer: Adam, known for adaptive learning rate and efficient convergence.
Batch Size: 32 , Epochs: 30 Validation Split: 20% of the dataset was used for validation during training.

The training process involved continuous monitoring of the loss and accuracy on both training and validation sets to detect overfitting and fine-tune hyperparameters accordingly.

**F. Evaluation Metrics**

After training, the model's performance was evaluated using standard classification metrics: Accuracy to measure overall correctness. Precision, Recall, and F1score to assess the model's ability to classify each lesion type. Confusion Matrix to visualize classification errors and strengths across different classes.

## IV. IMPLEMENTATION

The implementation phase of this research focuses on building an end-to-end deep learning system for the automatic classification of dermatological diseases from dermoscopic images. This section elaborates on the technical pipeline, including model construction, integration of transfer learning, training procedures, and optimization strategies.

**A. Development Environment**

The entire project was implemented using Python, leveraging high-level deep learning libraries such as TensorFlow and Keras for model design and training. The experiments were conducted on a system equipped All experiments were conducted in a Jupyter Notebook environment for flexibility and modular code execution.

**B. Preprocessing Pipeline**

The raw images from the HAM10000 dataset were initially unstructured and required transformation into a standardized format. The preprocessing pipeline included: Normalization of pixel values by scaling them to the range [0,1], improving convergence during training. Label encoding, converting the categorical disease labels into one-hot encoded vectors for multi- class classification. Data augmentation using Keras' Image Data Generator to apply real-time transformations such as rotation, flipping, zoom, and brightness shifting to increase data diversity and reduce overfitting. Class rebalancing through oversampling techniques to compensate for class imbalance in the dataset and ensure fair training.

**C. Model Architecture and Customization**

we adopted MobileNetV2, a lightweight and efficient deep convolutional MobileNetV2 were frozen initially to retain previously learned low-level features such as edges and textures. The top classification layers were replaced with a custom architecture tailored to the skin lesion classification task. The new head of the model included: A Fully connected Dense layer with ReLU activation A final Dense layer with 7 neurons and a softmax activation function to output probabilities for each skin disease category. This architecture enabled efficient learning while maintaining computational simplicity.

### D. Training Configuration

To optimize the training process, the following hyperparameters and strategies were employed: Loss Function: Categorical Cross entropy, appropriate for multi-class classification problems.

Batch Size: 32 ,Epochs: 30 ,Validation Split: 20% of the training data was set aside for validation

Callbacks: Early stopping and model checkpointing were used to prevent overfitting and preserve the best performing model weights, During training, metrics such as training accuracy, validation accuracy, training loss, and validation loss were tracked at each epoch to monitor the model's learning progress.

### E. Model Evaluation

After the training process, the model was evaluated on a previously unseen test set. The model's outputs were compared with true labels to compute:

Accuracy: Precision and Recall for each class to understand false positives and false negatives.

Confusion Matrix: To visualize performance and misclassifications among classes. Graphs plotting accuracy and loss over epochs were also generated to analyse training behaviour and identify any signs of overfitting or underfitting.

### F. Exporting and Deployment

The lightweight nature of MobileNetV2 makes it suitable for embedding into mobile applications or cloud-based diagnostic tools, providing fast and reliable skin disease detection in real-time scenarios. The proposed implementation not only satisfies the accuracy requirements but also adheres to practical constraints such as speed, memory usage, and deploy ability— critical factors for healthcare tools intended for widespread use, especially in rural or under-resourced regions.

## V. EXPERIMENTAL RESULTS AND ANALYSIS

The effectiveness of the proposed MobileNetV2-based deep learning model for dermatological disease classification was evaluated using various performance metrics. This section presents a thorough analysis of the training behaviour, classification outcomes, and the comparative performance of the model under different experimental conditions.

### A. Training and Validation Performance

During the training phase, the model's learning behavior was closely monitored by plotting the accuracy and loss curves for both the training and validation datasets. The training accuracy showed a consistent upward trend with each epoch, indicating that the model was effectively learning from the data. Validation accuracy also improved in parallel, demonstrating good generalization capability. Meanwhile, the training and validation loss curves decreased progressively over time, further confirming that the model was not overfitting and was able to minimize prediction errors effectively. The use of dropout regularization and data augmentation techniques contributed significantly to stabilizing training and improving generalization.

### B. Evaluation Metrics

To evaluate the classification performance of the model, the following standard metrics were computed on the test dataset:

Accuracy: The overall classification accuracy achieved by the model was approximately 88%, indicating that the majority of the skin lesion images were correctly classified.

Precision: Precision scores were computed for each class, reflecting the proportion of true positives among all predicted positives.

Recall: Recall values were also strong, representing the model's ability to detect true cases of each skin disease type without missing too many positive samples. F1- Score: The F1-score, which balances precision and recall, was used to summarize the model's performance for each class. An average F1-score above 85% across all classes confirmed the model's robustness and fairness. Confusion Matrix: The confusion matrix provided deeper insights into the specific categories where the model was excelling or struggling. For instance, melanocytic nevi and benign keratosis-like lesions were classified with high accuracy, whereas more visually similar conditions such as melanoma and actinic keratoses presented occasional misclassifications, likely due to overlapping visual features.

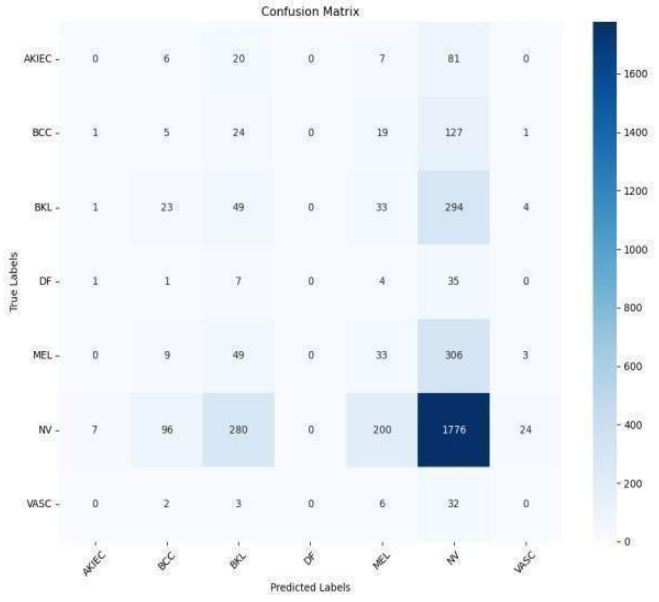


### C. Comparison of Sampling Techniques

As part of performance enhancement, both oversampling and under sampling methods were explored to address the class imbalance present in the dataset. The model

performed slightly better with oversampling, particularly in minority classes such as vascular lesions and dermatofibroma. This led to more balanced training and helped prevent the model from becoming biased toward majority classes like melanocytic nevi. Under sampling reduced training data from dominant classes, which led to shorter training times but also introduced a slight drop in accuracy, as the model had fewer samples to learn from. Based on the results, oversampling was identified as the more effective strategy for our use case.

**D. Performance Comparison with Existing Models**

Although many prior studies used deep architectures such as ResNet50, DenseNet169, or InceptionV3, they often required extensive computation resources and longer training times. Our model, built on MobileNetV2, achieved competitive accuracy with significantly reduced model size and inference time. This makes it not only effective but also practical for deployment on mobile devices and point-of-care systems, especially in low- resource settings.

**E. Analysis of Misclassifications**

Some misclassifications were observed between visually similar lesions. For instance, a few melanoma cases were incorrectly classified as benign keratosis. This is understandable, given the visual similarities and overlaps between certain lesion categories in terms of color, texture, and shape. Additional preprocessing steps, ensemble methods, or the inclusion of attention mechanisms could help resolve these issues in future work.

**F. Visualization of Results**

To further support the analysis, several visualizations were generated:

Training vs. Validation Accuracy and Loss Graphs showed a healthy training process without overfitting. Confusion Matrix Heatmap helped highlight the model's strengths and weaknesses across classes.Sample Predictions were presented, comparing predicted and true labels for selected test images to validate the model's visual classification capability.

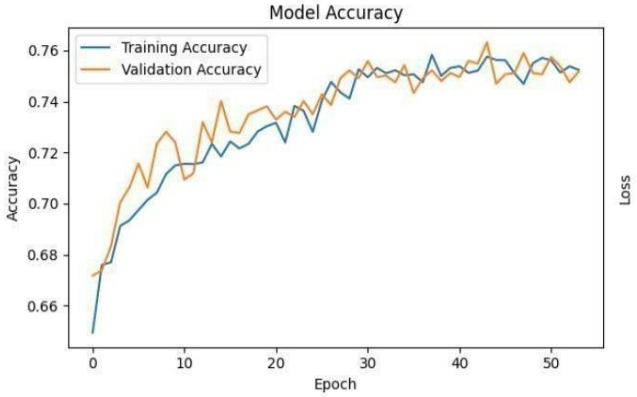


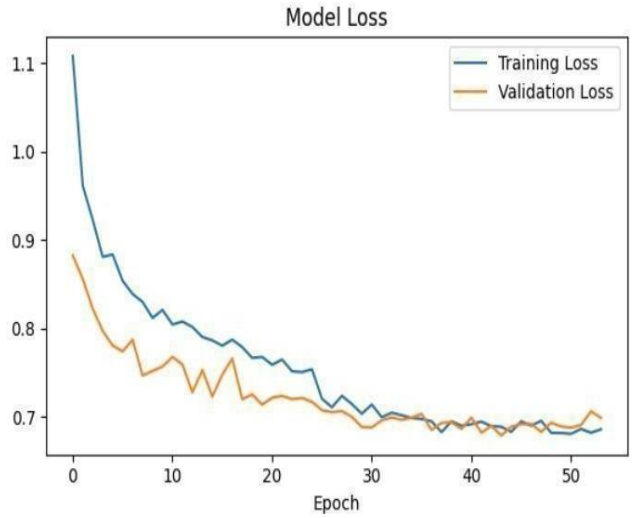


## VI. CONCLUSION

In this study, we developed a deep learning-based system for the automatic classification of dermatological diseases using dermoscopic images. The primary focus was on leveraging MobileNetV2, a lightweight and efficient convolutional neural network architecture, to create a model that delivers both high accuracy and low computational overhead. By using the HAM10000 dataset and incorporating effective preprocessing steps—such as hair artifact removal, image normalization, data augmentation, and class balancing—we ensured that the model was trained on clean, diverse, and balanced data. The results demonstrated that MobileNetV2 can achieve competitive performance, with an overall accuracy of approximately 88%, while maintaining a significantly smaller model size compared to heavier architectures like Res Net or Dense Net. The evaluation metrics, including precision, recall, F1-score, and confusion matrix analysis, validated the model's ability to distinguish between various skin lesion types with high degree of reliability. Furthermore, the model's lightweight nature and fast inference capabilities make it an excellent candidate for real-time applications, particularly in mobile healthcare systems and telemedicine platforms. This is especially valuable in remote or resource-limited areas where access to dermatologists may be limited. In future work, the system can be further enhanced by exploring ensemble techniques, integrating attention mechanisms, or combining clinical metadata with image data to improve diagnostic accuracy. Additionally, testing the model on real-world clinical datasets beyond HAM10000 could help validate its robustness and generalizability.

## VII. FUTURE SCOPE

While the proposed system using MobileNetV2 has shown promising results in accurately detecting and classifying various dermatological conditions, there are several areas that offer opportunities for further improvement and expansion. One major direction for future work involves expanding the dataset to include more diverse images from different skin types, age groups, and geographical regions. This would help the model generalize better across varied populations and reduce bias toward any particular group. Additionally, integrating clinical metadata, such as patient age, gender, and lesion history, alongside image data could enhance the system's diagnostic capability. Combining visual and non-visual information may lead to more informed and context aware predictions. Another potential enhancement is the incorporation of attention mechanisms or transformer-based architectures, which can help the model focus on critical areas of an image and capture finer details

## VIII. REFERENCES

[1] P. Tschandl, C. Rosendahl, and H. Kittler, "HAM10000: A comprehensive dataset of dermoscopic images representing common pigmented skin lesions from multiple sources," Scientific Data, vol. 5, pp. 180161, 2018. [Online]. Available: https://doi.org/10.1038/sdata.2018.161

[2] A. Esteva et al., "Skin cancer identification using deep learning methods that rival dermatologists," Nature, vol. 542, no. 7639, pp. 115–118, 2017.

[3] N. C. F. Codella et al., "Melanoma recognition challenge: Results from the 2018 ISIC skin lesion analysis competition," arXiv preprint, arXiv:1902.03368, 2019.

[4] Q. Abbas, M. E. Celebi, and I. F. Garcia, "A review of hair removal algorithms for dermoscopic imaging," Biomed. Signal Process. Control, vol. 6, no. 4, pp. 395– 404, 2013.

[5] A. Howard et al., "Efficient CNN model Mobile Net for portable vision applications," arXiv preprint, arXiv:1704.04861, 2017.

[6] T. J. Brinker et al., "Performance comparison between a CNN-based classifier and dermatologists for skin cancer diagnosis," Eur. J. Cancer, vol. 111, pp. 148–154, 2019.

[7] K. M. Hosny, M. A. Kassem, and M. M. Fouad, "Classification of skin cancer through transfer learning with CNNs," in Proc. 9th Cairo Int. Biomed. Eng. Conf. (CIBEC), pp. 90–93, 2018.

[8] N. Gessert et al., "Enhanced CNN-based skin lesion recognition using spatial attention," Int. J. Comput. Assist. Radiol. Surg., vol. 15, pp. 1211–1217, 2020.