



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 Issue: V Month of publication: May 2025

DOI: <https://doi.org/10.22214/ijraset.2025.71340>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

Healthcare Innovation for Skin Disease Analysis Using ML and DL

Dhruv Malhotra¹, Dhruv Kuchhal²

Delhi Technological University, India

Abstract: Skin diseases represent a significant health-care burden worldwide, and rapid diagnosis is essential for effective treatment. Machine learning (ML) and deep learning (DL) are increasingly applied to automate and improve dermatological diagnosis. In this study, we present a novel skin disease classification pipeline that integrates multiple ML and DL techniques in a unified framework. First, dermoscopic images are preprocessed to remove artifacts (e.g., hair) using morphological filters and inpainting, and segmented using clustering and GrabCut algorithms to isolate the lesion region. To augment the training data and address class imbalance, we train a generative adversarial network (GAN) on a multi-class skin dataset. The refined dataset is then fed into dual deep feature extractors (ResNet50 and DenseNet121, pretrained on ImageNet) to obtain rich feature representations. These features are combined through an attention-based fusion strategy and passed to a hybrid classifier comprising gradient boosting models and a support vector machine. Key contributions include the pipeline of artifact removal and segmentation, GAN-based data enhancement, and an ensemble classification strategy with explainable-AI components for transparency. We evaluate the approach on a publicly available multi-category skin disease dataset (21 classes), achieving high diagnostic accuracy (>90%) on the test set. Results indicate that our integrated ML/DL system outperforms traditional methods, reduces time to diagnosis, and supports efficient telemedicine for enhanced patient care. Moreover, this approach sets a foundation for real-world deployment in healthcare environments by ensuring model interpretability and scalability. The success of our framework demonstrates the transformative impact AI can have on improving dermatological diagnostics and patient outcomes, especially in underserved and remote areas where access to specialists is limited.

Keywords: Artificial Intelligence; Deep Learning; Dermatology; Machine Learning; Skin Disease; Telemedicine.

I. INTRODUCTION

Skin diseases such as melanoma, eczema, and psoriasis affect a large portion of the population and can have serious health consequences if not diagnosed early. Dermatological diagnosis is typically based on visual examination by a specialist, but many regions lack sufficient experts, leading to delays and errors. AI-driven tools aim to bridge this gap by providing automated analysis of skin images. In particular, ML and DL have shown great promise in medical imaging.

Recent studies have applied a variety of ML and DL methods to skin lesion classification. Ahammed et al.

[1] introduced a segmentation-based ML pipeline and showed that proper preprocessing improves accuracy. Sun et al. [2] reviewed ML techniques in dermatology, noting that ensemble and hybrid models often outperform single classifiers. Johnson et al. [3] compared CNN-based deep learning to traditional ML on skin lesion data, finding that CNNs significantly improve feature representation. Similarly, Bandyopadhyay et al.

[4] and Kalaivani & Karpagavalli [5] demonstrated improved classification by combining multiple CNN architectures or fusing deep features with tree-based classifiers. Al-Dera & Othman [6] proposed a hybrid approach integrating image processing with ML, achieving robust skin disease diagnosis. These studies demonstrate the potential of advanced ML and DL for dermatology.

Despite these advances, several challenges remain. Dermoscopic images often contain artifacts (hair, bubbles, uneven lighting) that can confuse algorithms if not properly removed. Many existing methods rely on simple preprocessing or ignore such artifacts, reducing robustness. Public skin disease datasets also tend to be imbalanced: common conditions have thousands of examples while rare diseases have very few, which biases the models. Furthermore, few prior systems incorporate model explainability or are designed for seamless telemedicine integration. Explainable AI (XAI) techniques (e.g., saliency maps) are important for clinical trust, yet are rarely included in dermatology classifiers.

Moreover, current diagnostic tools often lack the flexibility to generalize across diverse skin types and environmental conditions. Skin disease appearance varies significantly with ethnicity, age, and lighting, demanding robust models that can adapt to these variations.

Traditional algorithms may also suffer from high false-positive or false-negative rates when exposed to unseen data. Incorporating cross-domain transfer learning and continual model updates is therefore crucial for maintaining high diagnostic performance over time.

To address these gaps, we propose a comprehensive multi-stage skin disease analysis framework. The pipeline begins with advanced preprocessing: a black-hat morphological filter and inpainting are applied to remove hair and noise, followed by k-means clustering and the GrabCut algorithm to precisely segment the lesion. Next, we train a generative adversarial network (GAN) on the available images to generate realistic synthetic lesion examples, augmenting underrepresented classes. We then extract deep features using two pretrained CNNs (ResNet50 and DenseNet121) and fuse them through an attention mechanism. Finally, a hybrid classifier (an ensemble of gradient-boosting trees and an SVM) is trained for diagnosis, and Grad-CAM is used to provide visual explanations of model decisions.

The remainder of the paper is organized as follows: Section II describes the dataset and methodology in detail. Section III presents the experimental results and evaluation metrics. Section IV discusses the implications of our findings, limitations of the current approach, and future work. Section V concludes with a summary of our contributions to healthcare innovation in dermatology.

II. MATERIALS AND METHODS

A. Dataset

We used a publicly available dermoscopic dataset containing thousands of images across multiple skin disease categories. The training set comprised 6,450 images and the test set 3,521 images, covering eight disease categories (e.g., acne/rosacea, atopic dermatitis, eczema, warts, fungal infections, etc.). Each image was resized to 224×224 pixels and normalized to $[0, 1]$ range to match the input requirements of the CNNs. Data split follows a fixed train/test partition provided by the dataset, ensuring no overlap between sets.

To improve the robustness of our models, we applied additional data augmentation techniques such as horizontal flipping, random rotations, brightness adjustments, and zooming. This helped simulate real-world variability and improve generalization. The dataset was further analyzed for intra-class variance to identify classes that were visually similar, thus guiding the design of discriminative feature extraction strategies.

B. Preprocessing

Artifact Removal: To enhance image quality, we first removed hair and noise artifacts. A black-hat morphological filter (the difference between the closing of the image and the original) was applied to highlight hair-like structures on the skin. The detected hair pixels were then removed using an inpainting algorithm, which fills in these regions based on surrounding pixel information. This produces smooth, hair-free images without altering the lesion.

Additionally, uneven lighting artifacts were mitigated using histogram equalization in the HSV color space. By adjusting image contrast uniformly, this step ensures that lesion features remain prominent and detectable by the subsequent segmentation and classification modules.

Lesion Segmentation: After artifact removal, we isolated the lesion region. The image was converted to HSV color space for improved segmentation. We applied k-means clustering (with $k = 3$) on the pixel colors to partition the image; one cluster typically corresponds to the lesion. The initial mask from k-means was refined using the GrabCut algorithm, which formulates segmentation as an energy minimization problem combining color statistics and a smoothness prior. The final binary mask precisely outlines the lesion, which is then extracted for feature analysis.

To validate segmentation quality, we compared the extracted masks visually with expert annotations available for a subset of images. A satisfactory overlap was observed, suggesting that the segmentation step preserved essential lesion information critical for accurate classification.

C. Data Augmentation (GAN)

We implemented a generative adversarial network (GAN) to synthesize realistic skin lesion images for data augmentation. The GAN consists of a generator G and a discriminator D trained in an adversarial manner:

G maps a random noise vector $\mathbf{z} \sim N(0, I)$ (optionally conditioned on class label) to an image $G(\mathbf{z})$, while D

attempts to distinguish $G(\mathbf{z})$ from real images. We used a convolutional architecture for both networks (similar to DCGAN).

The GAN was trained for 2000 epochs using the binary cross-entropy loss:

$$L = -E_{x \sim p_{data}}[\log D(x)] - E_{z \sim p_z}[\log(1 - D(G(z)))].$$

After training, the generator can produce new lesion images for each disease class, augmenting the dataset. In practice, we generated approximately one synthetic image for each original, effectively doubling the training data size and balancing the class distribution.

Furthermore, we employed a filtering mechanism to discard low-quality synthetic images by calculating their Frechet Inception Distance (FID) scores against real images, ensuring only high-fidelity samples were used for training.

D. Feature Extraction and Classification

For feature extraction, we used two deep CNNs pretrained on ImageNet: ResNet50 and DenseNet121. We removed the top classification layers and applied global average pooling to the final convolutional feature maps, resulting in 2048-dimensional (ResNet) and 1024-dimensional (DenseNet) feature vectors. These vectors were concatenated into a 3072-dimensional feature descriptor for each image. We also passed this concatenated vector through a learned attention module to weight the contributions of each network's features before classification.

The fused feature vectors were used to train multiple classifiers:

- Support Vector Machine (RBF): A one-vs-rest SVM with radial basis function kernel.
- Random Forest: An ensemble of 500 decision trees.
- XGBoost: A gradient-boosted decision tree model with 100 estimators.
- LightGBM: A gradient-boosted tree model (100 estimators) optimized for speed.

In addition, we built a stacked ensemble: the prediction probabilities from the above models were used as input features for a meta-classifier (logistic regression), which outputs the final class. All classifiers were trained on the fused CNN features from the training set. We used the Adam optimizer (learning rate 1×10^{-4}) and categorical cross-entropy loss for training the CNN

feature extractor networks. Classical ML models were trained using standard libraries (scikit-learn, XGBoost) with default hyperparameters unless noted.

TABLE I
SUMMARY OF MODELS AND TECHNIQUES.

Model/Technique	Description
ResNet50	50-layer CNN with residual connections (ImageNet-pretrained)
DenseNet121	121-layer densely-connected CNN (ImageNet-pretrained)
Support Vector Machine (RBF)	Kernel-based classifier (RBF kernel)
Random Forest	Ensemble of decision trees (500 trees)
XGBoost	Gradient boosting decision trees (100 estimators)
LightGBM	Gradient boosting with leaf-wise splitting (100 estimators)

III. RESULTS

We evaluated the models using standard metrics: accuracy, precision, recall, F1 score, and confusion matrix. Results are reported on the test set comprising unseen images. Our hybrid system achieved an overall test accuracy of 92.3%.

A. Quantitative Evaluation

Table 2 shows the classification performance across models. The stacked ensemble achieved the highest accuracy and F1 score, outperforming individual classifiers.

TABLE II
PERFORMANCE OF DIFFERENT CLASSIFIERS ON TEST DATA.

Model	Accuracy	Precision	Recall	F1-score
SVM (RBF)	88.1%	88.3%	87.9%	88.0%
Random Forest	89.5%	89.7%	89.2%	89.4%
XGBoost	90.6%	90.8%	90.4%	90.6%
LightGBM	90.9%	91.0%	90.7%	90.8%
Stacked Ensemble	92.3%	92.5%	92.1%	92.3%

B. Confusion Matrix and ROC Analysis

The confusion matrix (Figure 1) shows that most classes are well-classified, with few misclassifications between visually similar diseases (e.g., eczema vs atopic dermatitis). ROC curves (Figure-2) confirm strong discriminative performance across all disease categories, with area-under-curve (AUC) values above 0.95 for all classes.




Fig.1. Confusion matrix of stacked ensemble classifier.




Fig.2. ROC curves per disease category.

C. Explainability Results

To improve model trust, we visualized Grad-CAM saliency maps, highlighting regions used by the classifier for each decision. As seen, the model focuses correctly on lesion regions rather than irrelevant background, suggesting reliable attention mechanisms. Such explainability tools are essential for clinical deployment to ensure doctors can verify model decisions.

IV. DISCUSSION

Our results demonstrate that a comprehensive pipeline combining preprocessing, GAN augmentation, feature fusion, and hybrid classification significantly enhances skin disease classification accuracy compared to standalone CNNs or classical ML models. Notably, the use of GANs to synthesize rare classes proved crucial for balancing the dataset and improving recall on underrepresented diseases.

Moreover, ensemble methods, especially stacking diverse models, consistently boosted classification metrics, confirming prior findings in ML that model diversity leads to better generalization. The combination of deep features extracted from ResNet and DenseNet captures complementary information: ResNet captures global contextual information via residual mappings, while DenseNet captures fine-grained patterns via dense connections. Attention-based fusion further improves feature weighting, ensuring the model focuses on the most informative representations.

The success of our explainable AI component (Grad-CAM) indicates that the classifier is not simply overfitting to noise but instead learns to localize disease regions. This improves clinical trust in real-world applications.

However, certain limitations remain. Despite augmentation, rare diseases with highly variable presentations (e.g., fungal infections) occasionally suffer misclassification. Moreover, while Grad-CAM offers coarse localization, it lacks fine-grained interpretability. Future work could explore more advanced XAI methods like LIME or SHAP for dermatology.

Another limitation involves demographic diversity: although the dataset includes various skin tones, it may not fully capture the global diversity found in real-world populations. Future datasets and models must ensure equitable performance across different ethnic backgrounds to prevent biases in automated diagnosis.

Additionally, deployment considerations such as model size and inference speed are important for telemedicine applications, especially in resource-constrained environments. Exploring lightweight CNN architectures (e.g., MobileNet) and model quantization could make the system more feasible for mobile deployment.

Finally, incorporating patient metadata (age, symptoms, location of lesion) alongside images could enhance model context-awareness and improve diagnostic specificity.

V. CONCLUSION

We proposed a novel integrated ML/DL framework for automated skin disease diagnosis, combining advanced artifact removal, lesion segmentation, GAN-based data augmentation, dual CNN feature extraction with attention fusion, and hybrid ensemble classification. Extensive evaluation shows that our system achieves state-of-the-art diagnostic accuracy (>92%) across multiple skin diseases. Explainable AI techniques (Grad-CAM) confirm model transparency, critical for clinical adoption.

This research demonstrates that combining preprocessing, data augmentation, diverse feature extraction, and ensemble strategies can substantially improve dermatological AI systems. The proposed pipeline holds promise for real-world teledermatology applications, providing accessible and accurate skin disease diagnosis to underserved populations. Future work will extend this approach to broader datasets, incorporate metadata, and optimize the system for deployment on mobile and edge devices.

REFERENCES

- [1] Ahammed, S., et al. (2020). "Skin lesion classification using machine learning with preprocessing." *Biomedical Signal Processing and Control*, 57, 101819.
- [2] Sun, X., et al. (2021). "Machine learning and deep learning for skin disease classification: A review." *Journal of Healthcare Engineering*, 2021, 1-19.
- [3] Johnson, K.W., et al. (2022). "Comparing deep learning to traditional machine learning for skin lesion diagnosis." *Artificial Intelligence in Medicine*, 124, 102184.
- [4] Bandyopadhyay, S., & Dutta, S. (2018). "Fusion of CNN features for skin disease classification." *International Conference on Computer Vision Systems*.
- [5] Kalaivani, P., & Karpagavalli, S. (2020). "Hybrid deep learning approach for skin lesion classification." *Materials Today: Proceedings*, 33, 4863-4870.
- [6] AlDera, R., & Othman, A. (2018). "A hybrid approach for skin disease detection using image processing and machine learning." *International Journal of Advanced Computer Science and Applications*, 9(12), 146-151.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)