

Skin Disease Type Detector Using Deep Learning

PARTH SHARMA¹, NILESH KUMAR PANDEY², DR. ISHRAT ALI³, DR. SANJAY PACHAURI⁴
^{1,2,3,4}*Department of Data Science (DDCS), GNIOT College, Greater Noida, India*

Abstract- The exponential rise in skin diseases across the globe has created an urgent demand for automated, scalable, and precise diagnostic tools. In this paper, we present a deep learning-based Skin Disease Type Detector that leverages Convolutional Neural Networks (CNNs) and transfer learning to classify multiple types of skin conditions from dermoscopic and clinical images. The system is built on a modular image processing pipeline that pre-processes skin lesion images, extracts spatial and textural features through deep neural layers, and produces accurate multi-class predictions with confidence scores. Our implementation integrates a user-friendly interface that allows clinicians and patients to upload images and receive instant, cited predictions grounded in medical image databases. Experimental results demonstrate significant improvements in classification accuracy when compared to conventional machine learning baselines, while also reducing misdiagnosis rates. The system is designed to be domain-adaptive and can be extended to accommodate new disease categories through transfer learning without retraining from scratch. This work represents a meaningful step toward building trustworthy AI systems that combine generative fluency with clinical reliability for dermatological diagnosis.

Index Terms- Skin Disease Detection, Deep Learning, Convolutional Neural Networks, Transfer Learning, Image Classification, Dermoscopy, Medical AI, HAM10000, ResNet, EfficientNet.

I. INTRODUCTION

The prevalence of skin diseases worldwide has reached alarming proportions, with the World Health Organization estimating that skin conditions affect nearly 900 million people at any given time. These range from benign conditions such as eczema and psoriasis to life-threatening malignancies like melanoma. Despite the scale of the problem, accurate early diagnosis remains a significant challenge, especially in regions with limited access to specialist dermatologists. Traditional diagnosis relies heavily on clinical examination and biopsy, processes that are

time-consuming, expensive, and prone to human error.

Recent advances in deep learning and computer vision have opened promising avenues for automated skin disease detection. Convolutional Neural Networks (CNNs) have demonstrated remarkable ability to identify patterns in medical images, often matching or exceeding the performance of trained clinicians in controlled evaluations. However, deploying such systems in real-world healthcare settings presents unique challenges, including the need for diverse training data, robustness to image variability, and interpretability for clinical trust.

This paper introduces a comprehensive implementation of a Skin Disease Type Detector, which leverages state-of-the-art CNN architectures and transfer learning to classify images into multiple disease categories. Our system addresses key limitations of prior work by employing a modular pipeline that supports real-time inference, confidence-scored predictions, and transparent source attribution. The key contributions of this work include: (1) a robust image pre-processing and augmentation pipeline that improves model generalization across diverse image sources, (2) implementation of pre-trained deep learning architectures fine-tuned for skin disease classification, (3) development of a user-friendly diagnostic interface providing disease category predictions with confidence scores, and (4) empirical evaluation demonstrating high classification accuracy and improved sensitivity across multiple disease types.

II. LITERATURE REVIEW

The application of machine learning to dermatological image analysis has evolved substantially over the past two decades. This section

surveys the foundational and recent work that informs our approach.

A. Historical Context

Early automated skin lesion analysis systems relied on hand-crafted feature extraction techniques such as ABCD rule (Asymmetry, Border, Color, Diameter) analysis and texture descriptors like Local Binary Patterns (LBP). These systems performed adequately on controlled datasets but failed to generalize across varying imaging conditions and patient demographics. The advent of machine learning, particularly Support Vector Machines (SVMs) and Random Forests, improved classification accuracy but still depended on manually engineered features, limiting scalability and robustness.

B. Practical Frameworks and Toolkits

The introduction of ImageNet and large-scale visual recognition challenges catalyzed the adoption of deep CNNs in medical imaging. Frameworks such as TensorFlow and PyTorch have since made it practical to prototype and deploy deep learning models for skin disease classification. Pre-trained models such as VGG, ResNet, Inception, and EfficientNet have been fine-tuned on publicly available dermatological datasets like HAM10000, ISIC, and PH2, enabling researchers and developers to leverage transfer learning for rapid and effective model development.

C. Domain-Specific Applications

Deep learning-based skin disease detection systems have been deployed across multiple clinical contexts. Esteva et al. demonstrated that CNNs could classify skin cancer with dermatologist-level accuracy. In telemedicine, systems have been developed for remote triage of skin conditions using smartphone images. Explainability tools such as Grad-CAM have been incorporated to highlight lesion regions that influence model predictions, increasing clinician trust. Our work builds upon these advancements to create an end-to-end diagnostic tool that is accessible, accurate, and interpretable.

III. METHODOLOGY

A. Image pre-processing

The system accepts skin lesion images in common formats including JPEG, PNG, and TIFF. Upon

upload, images undergo a standardized pre-processing pipeline: resizing to 224x224 pixels, normalization of pixel intensities to the [0,1] range, and color space conversion where necessary. To improve model robustness, data augmentation techniques including random horizontal and vertical flips, rotation up to 30 degrees, brightness and contrast jitter, and zoom augmentation are applied during training.

B. Feature Extraction and Model Architecture

Feature extraction is performed using pre-trained CNN architectures, primarily EfficientNet-B4 and ResNet-50, both fine-tuned on the HAM10000 dermatology dataset containing over 10,000 labeled dermoscopic images across seven disease categories. The final classification layers are replaced with custom dense layers and a softmax output corresponding to the target disease classes. Dropout regularization is applied to prevent overfitting.

C. Classification and Prediction

The classification module receives deep feature vectors from the CNN backbone and produces a probability distribution over disease categories using a fully connected softmax layer. Predictions are accompanied by confidence scores, allowing users to assess reliability. When confidence falls below a defined threshold, the system prompts the user to consult a medical professional, reinforcing responsible AI use.

IV. RESULTS

A. Performance Analysis

Experimental results demonstrate consistent and significant improvements of the proposed Skin Disease Type Detector over baseline machine learning models. The deep learning system outperformed SVM and Random Forest baselines across all evaluated disease categories in terms of classification accuracy, sensitivity, and specificity. The EfficientNet-B4 model achieved an overall accuracy of 91.3% on the HAM10000 test set, compared to 74.8% for the SVM baseline. Predictions produced by the system were directly verifiable against labeled ground truth images in the evaluation database.

B. Effects

The most significant benefit of the deep learning architecture was its ability to distinguish visually similar disease classes that are frequently confused in clinical practice, such as melanoma and benign melanocytic nevi. The attention mechanisms embedded within EfficientNet enabled the model to focus on clinically relevant features such as irregular borders and heterogeneous pigmentation, resulting in substantially fewer false negative classifications across every evaluation category.

C. Limitations

While effective, several limitations were identified. First, system performance is highly dependent on image quality and consistency of capture conditions. Images taken under poor lighting or with significant motion blur resulted in degraded classification accuracy. However, the confidence score mechanism provides transparency, prompting users when predictions may be unreliable. Second, the current dataset predominantly represents light-skinned individuals, and performance was observed to be lower on darker skin tones, highlighting the need for more diverse and representative training data in future iterations.

V. CONCLUSION

This research demonstrates that deep learning-based skin disease detection significantly enhances diagnostic accuracy and efficiency in dermatological AI. The proposed Skin Disease Type Detector achieved state-of-the-art classification performance across multiple disease categories while maintaining interpretability through confidence scoring and attention visualization. Future work may include expanding the training dataset to cover a broader range of skin tones and disease subtypes, integrating multimodal inputs such as patient history and dermoscopy metadata, and incorporating federated learning strategies to enable privacy-preserving model updates across hospital systems.

REFERENCES

[1] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, pp. 115-118, 2017.

[2] N. C. F. Codella et al., "Skin lesion analysis toward melanoma detection: ISIC 2018 challenge," in *Proc. ISBI*, 2019.

[3] P. Tschandl et al., "The HAM10000 dataset, a large collection of multi-source dermoscopic images of common pigmented skin lesions," *Scientific Data*, vol. 5, 2018.

[4] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in *Proc. ICML*, 2019, pp. 6105-6114.

[5] K. He et al., "Deep residual learning for image recognition," in *Proc. CVPR*, 2016, pp. 770-778.

[6] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proc. ICCV*, 2017, pp. 618-626.

[7] J. Kawahara et al., "Seven-point checklist and skin lesion classification using multitask multimodal neural nets," *IEEE Journal of Biomedical and Health Informatics*, vol. 23, no. 2, pp. 538-546, 2019.

[8] X. Li et al., "Skin lesion analysis towards melanoma detection using deep learning network," *Sensors*, vol. 18, no. 2, 2018.

[9] A. Krizhevsky et al., "ImageNet classification with deep convolutional neural networks," in *Proc. NeurIPS*, 2012, pp. 1097-1105.

[10] Y. LeCun et al., "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278-2324, 1998.

[11] "PyTorch Documentation," [Online]. Available: <https://pytorch.org/docs/>

[12] A. Mahbod et al., "Transfer learning using a multi-scale and multi-network ensemble for skin lesion classification," *Computer Methods and Programs in Biomedicine*, vol. 193, 2020.

[13] S. Haenssle et al., "Man against machine: Diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition," *Annals of Oncology*, vol. 29, no. 8, 2018.

[14] T. J. Brinker et al., "Skin cancer classification using convolutional neural networks," *Journal of Investigative Dermatology*, vol. 138, no. 11, 2018.

[15] G. Huang et al., "Densely connected convolutional networks," in *Proc. CVPR*, 2017, pp. 4700-4708.

[16] C. Szegedy et al., "Inception-v4, inception-ResNet and the impact of residual connections on learning," in *Proc. AAAI*, 2017.