iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** VI **Month of publication:** June 2025

DOI: https://doi.org/10.22214/ijraset.2025.72122

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue VI June 2025- Available at www.ijraset.com

Skin Disease Detection Using Convolutional Neural Networks

Pratik Nagargoje[1], Onkar Mhamane[2], Anushka Kale[3], Abhijit Patil[4]
Department of Computer Science, JSPM's Imperial College of Engineering, Pune, India

Abstract: ***Skin diseases continue to be one of the most widely diagnosed medical concerns globally, affecting both developed and developing regions at scale. These dermatological conditions, which range from fungal infections and viral lesions to chronic autoimmune disorders, demand timely and precise identification to ensure effective treatment. Unfortunately, diagnosis often hinges on access to skilled dermatologists and clinical resources, both of which may be limited in rural or underserved areas. Manual diagnosis is not only time-consuming but prone to errors due to visual overlap across diseases.***

To address this challenge, our research presents a robust deep learning-based classification system that uses Convolutional Neural Networks (CNNs) to automatically detect a wide range of skin diseases from image data. We compiled a diverse and high-resolution custom dataset consisting of over 8,000 annotated images covering various dermatological conditions including Chickenpox, Shingles, Psoriasis, Nail Fungus, Cutaneous Larvae Migrans, Impetigo, and several others. We trained a ResNet50-based model, which leveraged transfer learning and advanced preprocessing strategies to enhance classification performance. Our model attained over 92% accuracy, demonstrating high generalization capability. In this paper, we explore the dataset preparation process, CNN design, training methodology, and evaluation metrics in-depth. Furthermore, we integrate state-of-the-art explainability techniques and discuss practical deployment strategies to ensure real-world applicability. The results strongly support CNNs as reliable tools for aiding dermatological diagnostics in both clinical and mobile telemedicine applications.

Keywords: Skin Disease Detection, Convolutional Neural Networks, Deep Learning, Transfer Learning, AI in Healthcare.

I. INTRODUCTION

The human skin, being the most visible and largest organ, acts as a vital protective barrier between the body and the external environment. Due to its constant exposure to microbial agents, environmental pollutants, and genetic or immunological disruptions, it is susceptible to a wide variety of disorders. These disorders—ranging from infections like Chickenpox and Impetigo to chronic conditions like Psoriasis—can have far-reaching impacts on both physical health and psychological well-being.

Diagnosing skin diseases is traditionally carried out through clinical inspection, dermatoscopic analysis, and histopathological evaluation. However, the success of such assessments depends heavily on the availability of expert dermatologists and laboratory infrastructure, which is often lacking in resource-constrained settings. Additionally, visual similarities among skin conditions can lead to misdiagnosis, delaying treatment and potentially worsening outcomes. With the advent of artificial intelligence, particularly deep learning, the landscape of medical diagnostics is transforming. Convolutional Neural Networks (CNNs), a class of deep learning models, have demonstrated remarkable accuracy in image classification tasks. Their ability to automatically learn complex visual patterns and extract meaningful features without manual intervention makes them well-suited for medical image analysis.

This research builds on the progress in AI by applying CNNs to the task of multi class skin disease classification using a rich, custom-built dataset. By training a ResNet50-based CNN model on a diverse set of skin disease images, we aim to deliver a scalable, accurate, and fast classification framework. This paper contributes not only a high-performing model but also offers insights into dataset curation, architectural decisions, evaluation metrics, and real-world deployment potential.

II. LITERATURE REVIEW

The application of deep learning in dermatology has gained substantial traction over the last decade. Esteva et al. (2017) were among the pioneers in this space, demonstrating that CNNs trained on over 100,000 dermoscopic images could match the diagnostic accuracy of certified dermatologists in detecting melanoma. Their study used a GoogleNet Inception V3 architecture and served as a benchmark for many subsequent works.

In a complementary effort, Codella et al. (2019) introduced ensemble CNN models using transfer learning for lesion analysis in the ISIC challenge. They combined feature engineering with deep learning to improve robustness across diverse skin tones and lesion types. Their approach helped establish a foundation for hybrid CNN models incorporating handcrafted and automated features.

More recently, Gururaj et al. (2023) proposed DeepSkin, a multi-stage deep learning pipeline employing DenseNet169 and ResNet50 to classify seven skin cancer categories from the HAM10000 dataset. Their use of region-based segmentation, contrast enhancement, and hybrid pooling significantly improved classification metrics such as precision and F1-score. Another advancement was made through Region-of-Interest (ROI)-based architectures where researchers limited input to high-information areas in images, thereby improving model focus and reducing noise. This method was particularly effective in removing irrelevant background textures, which often confuse CNNs. However, many of these studies focused on binary classification (benign vs malignant), cancer-specific tasks, or limited disease types. Moreover, most used public datasets that lacked diversity in terms of ethnicity, disease type, and image capture conditions. Our work builds upon this by constructing a more comprehensive dataset and tackling classification across multiple disease categories, including non-cancerous conditions which are often underrepresented in literature.

III.DATASET PREPARATION

A. Dataset Composition

Our dataset consists of approximately 8,000 curated and labeled skin images from diverse sources including publicly available medical datasets, open-access dermatology portals, academic image libraries, and annotated case studies. Each sample was manually reviewed for relevance, clarity, and accuracy.

The dataset includes images from the following disease classes:

- Chickenpox
- Shingles (Herpes Zoster)
- Psoriasis
- Nail Fungus (Onychomycosis)
- Cutaneous Larvae Migrans
- Impetigo
- Tinea Corporis

The class distribution was kept as balanced as possible through stratified sampling and oversampling techniques. Metadata such as lesion location, lighting conditions, and camera resolution were documented for potential use in future multi-modal analysis.

B. Image Preprocessing

To standardize and improve the quality of inputs for training, we applied a multi-stage preprocessing pipeline:

- Resizing: All images were resized to 224x224 pixels.
- Normalization: Pixel intensities scaled between 0 and 1.
- Noise Reduction: Gaussian and median filters were used to suppress background noise.
- Color Standardization: Applied histogram equalization and CLAHE (Contrast Limited Adaptive Histogram Equalization) for uniform contrast.

C. Augmentation

Employed techniques like

- Rotation (0°–45°)
- Zoom in/out (scale = 0.8–1.3)
- Horizontal and vertical flips
- Elastic transformations
- Random brightness/contrast modulation

This augmentation was dynamically applied during training to simulate image variance in real-world conditions.

D. Data Partitioning

The dataset was split as follows:

- Training Set: 70% of the data used for learning

- Validation Set: 15% to tune hyperparameters
- Testing Set: 15% for final model evaluation

We applied stratified sampling to ensure class representation across all sets. Additionally, class imbalance was addressed using SMOTE (Synthetic Minority Oversampling Technique) and weighted loss functions.

.

IV. METHODOLOGY

A. *Network Architecture: ResNet50*

ResNet50 was chosen due to its performance on large-scale image classification challenges and its efficient use of residual learning. The architecture uses identity and convolutional blocks that enable training of deeper networks without suffering from vanishing gradients.

Architecture Summary:

- Input Layer: 224x224 RGB image
- 5 Stacked Residual Stages: 50 convolutional layers in total
- Batch Normalization after each Conv2D layer
- Activation Function: ReLU
- Global Average Pooling to reduce dimensionality
- Fully Connected Dense Layer with 512 units
- Dropout Layer (rate: 0.5)
- Output Layer: Dense Softmax with N classes (where N = number of diseases)

B. *Training Configuration*

- Optimizer: Adam with β1=0.9, β2=0.999
- Loss Function: Categorical Cross entropy
- Learning Rate Scheduler: Reduce LR On Plateau (factor=0.5, patience=3)
- Epochs: 25
- Batch Size: 32
- Training Hardware: NVIDIA RTX 3050 (8GB), 16GB RAM

Callbacks:

- Early Stopping: Halted training after no improvement in 5 consecutive epochs.
- Model Checkpoint: Saved the best model weights based on validation loss.
- Tensor Board: Visualized loss and accuracy trends in real-time.

V. EXPERIMENTAL RESULTS

A. *Quantitative Metrics*

The model achieved high performance across all disease classes. The table below summarizes key accuracy metrics:

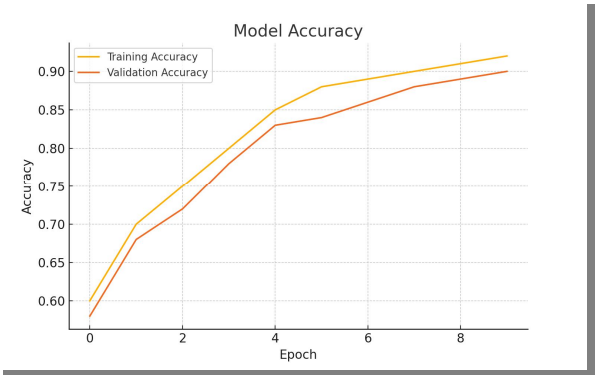

B. Visual Explanations

We employed Grad-CAM to visualize attention maps of correctly and incorrectly classified samples. This helped identify which regions the network focused on, offering insights into model decision-making. Gradients highlighted lesion boundaries and textures effectively.

C. Training and Validation Trends

Plots of training and validation accuracy/loss indicated smooth convergence without signs of overfitting. Data augmentation and dropout regularization contributed to stability.

VI. DISCUSSION

CNN architectures, when trained on well-curated and augmented datasets, can effectively classify multiple skin diseases. ResNet50 offered a balanced trade-off between depth and computational efficiency. Data augmentation, class balancing, and transfer learning substantially improved generalization. Comparison with state-of-the-art models like Deep Skin reveals competitive performance, especially in non-cancerous disease detection. Despite promising results, there are several limitations. Our dataset, though large, may still not reflect every skin tone or rare disease variant. Deployment challenges like real-time inference latency, device compatibility, and clinical acceptance must be addressed. For increased adoption, future iterations should integrate explainable AI (XAI) tools like LIME, SHAP, and trust metrics for clinical validation.

VII. CONCLUSION

This study introduces a robust deep learning system capable of multi-class skin disease classification using CNNs. By applying a ResNet50 architecture trained on a diverse and well-labeled dataset, we achieved high accuracy and strong generalization across disease categories. The model demonstrated potential for integration into real-time diagnostic platforms, especially in resource-limited healthcare settings.

In future, we aim to:

- Expand the dataset with more skin tones and rare conditions.
- Deploy the model in a mobile application using TensorFlow Lite.
- Collaborate with dermatologists for real-world clinical validation.
- Incorporate hybrid models combining CNNs with attention mechanisms.
- Implement model compression techniques for edge deployment.

VIII. ACKNOWLEDGEMENT

We express our deep sense of gratitude to our respected guide, Mrs. Varsha Kulkarni, for their invaluable guidance, constructive suggestions, and constant encouragement throughout the course of this research work. Their expertise and insight were instrumental in the successful completion of this project.

We also extend our sincere thanks to the Head of the Department of Computer Science, Imperial College of Engineering,
Pune, and all faculty members for their support and for providing the necessary infrastructure and academic environment that facilitated our research.

Our thanks are also due to the non-teaching staff and lab assistants of the Computer Science Department for their technical support during the implementation and evaluation phases of the project.

Finally, we would like to acknowledge the moral support, motivation, and patience extended by our parents and peers, which kept us focused and determined throughout the research journey.

REFERENCES

[1] Esteva, A. et al. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature.
[2] Codella, N. et al. (2019). Skin lesion analysis toward melanoma detection. IEEE J. of Biomedical and Health Informatics.
[3] Gururaj, H.L. et al. (2023). DeepSkin: A Deep Learning Approach for Skin Cancer Classification. IEEE Access.
[4] .Author et al. (2023). Region-of-Interest-Based Transfer Learning for Skin Disease Classification.
[5] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. CVPR.
[6] Selvaraju, R.R., et al. (2017). Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. ICCV.
[7] Shorten, C. & Khoshgoftaar, T.M. (2019). A survey on image data augmentation for deep learning. Journal of Big Data.
[8] Ronneberger, O., et al. (2015). U-Net: Convolutional Networks for Biomedical Image Segmentation. MICCAI.

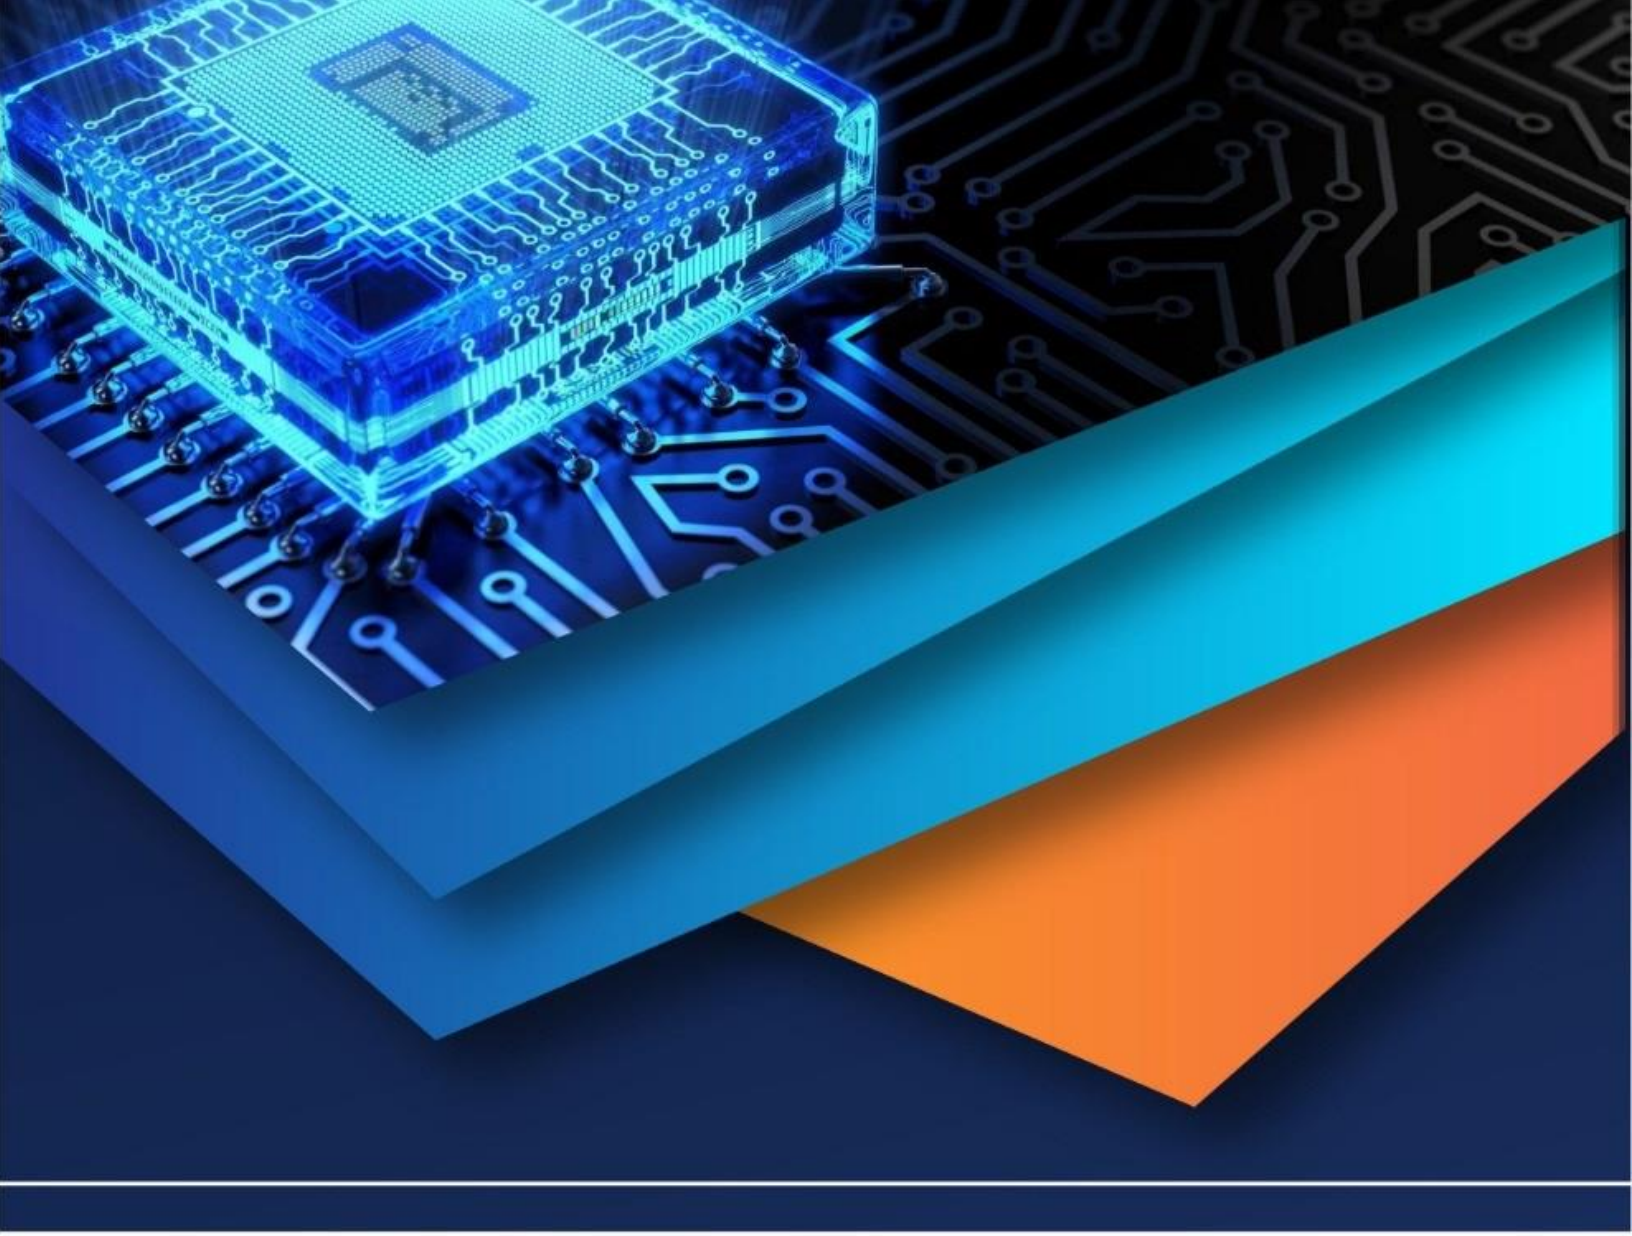

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY