Indian Journal of Computer Science and Technology
Volume 5, Issue1 (January-April 2026), PP: 413-416.
www.indjcst.com ISSN No: 2583-5300

Design and Implementation of a CNN-Based Web Application for Skin Disease Detection

G.Divya*[1], K.Bharathi*[2], T.Charulatha*[3], I.Harika*[4], G.Jothika*[5]
[1]*Assistant Professor, Department of Information Technology, PSV College of Engineering and Technology, Krishnagiri, Tamil Nadu, India.*
[2,3,4,5] *UG Scholars, Department of Information Technology, PSV College of Engineering and Technology, Krishnagiri, Tamil Nadu, India.*

To Cite this Article: *G.Divya*[1], K.Bharathi*[2], T.Charulatha*[3], I.Harika*[4], G.Jothika*[5], "Design and Implementation of a CNN-Based Web Application for Skin Disease Detection", Indian Journal of Computer Science and Technology, Volume 05, Issue 01 (January-April 2026), PP: 413-416.*

Abstract: *Skin diseases are among the most prevalent medical conditions worldwide, affecting millions of people across all age groups. Early and accurate diagnosis of skin conditions is essential to prevent serious health consequences, including skin cancer. However, dermatological expertise is limited in many regions, making access to timely diagnosis a significant challenge. This project presents DermaAI, an AI-powered skin disease classification system that leverages deep learning techniques to automatically identify and analyze common skin conditions from uploaded medical images The proposed system uses EfficientNet-B0 with transfer learning as the core deep learning architecture. The model is trained on a comprehensive dataset combining the DermNet NZ Image Library and the ISIC 2019 Challenge Dataset, enabling it to classify 23 distinct skin conditions including acne, melanoma, psoriasis, eczema, basal cell carcinoma, and tinea. The system incorporates a robust preprocessing pipeline consisting of image resizing to 224x224 pixels, normalization using ImageNet mean and standard deviation, and extensive data augmentation techniques including rotation, horizontal flipping, color jitter, and brightness adjustment.The DermaAI system is deployed as a Flask-based web application providing an intuitive drag-and-drop image upload interface with real-time prediction results, confidence scores, and detailed medical information. The system achieves an overall classification accuracy of 95.6% on the test dataset with an average prediction time of less than 2 seconds per image.*

Additional features include skin detection validation using HSV color space analysis, top-3 prediction display, treatment recommendations, and prominent medical disclaimers encouraging professional consultation.Experimental results demonstrate that the proposed system provides highly accurate, fast, and accessible dermatological screening support. The system has strong potential to improve healthcare accessibility, particularly in regions with limited dermatological services, while contributing to early detection of serious skin conditions.

I. INTRODUCTION

Skin is the largest organ of the human body and acts as the first line of defense against environmental factors, pathogens, and physical damage. Skin diseases range from mild, self- resolving conditions to life-threatening malignancies such as melanoma. According to the World Health Organization, skin diseases affect nearly one-third of the global population, making them one of the most widespread categories of human illness. Early detection and accurate diagnosis of skin conditions are critical for successful treatment and improved patient outcomes.

Traditional dermatological diagnosis relies heavily on the expertise and experience of trained dermatologists. However, the global shortage of dermatologists and uneven distribution of healthcare resources means that many patients, particularly in rural and developing regions, lack access to timely specialist care. Computer vision and deep learning have emerged as powerful tools in medical image analysis, offering the potential to automate and democratize dermatological screening.

This project presents DermaAI, an AI-powered skin disease analysis system that uses a deep learning model trained on thousands of dermatological images to classify 23 skin conditions with high accuracy. The system is deployed as a web application providing accessible, real-time skin analysis to users anywhere in the world.

Artificial intelligence and deep learning have demonstrated remarkable capabilities in medical image analysis, often achieving performance comparable to that of board-certified specialists. In dermatology, convolutional neural networks (CNNs) have proven particularly effective because skin disease diagnosis is inherently visual and pattern-based. AI-based systems can process large volumes of images consistently without fatigue, identify subtle patterns invisible to the naked eye, and provide quantitative confidence measures that aid clinical decision-making.

The development of large, publicly available dermatological image datasets such as the ISIC dataset and DermNet has facilitated the training of highly accurate deep learning models. Transfer learning, where models pre-trained on large natural image datasets are fine-tuned for medical tasks, has further accelerated progress by reducing the volume of labelled medical data required. These advances make AI-powered dermatological tools both technically feasible and practically valuable.

Skin disease diagnosis presents several challenges in current healthcare settings. First, dermatologist availability is severely limited in many parts of the world, creating long waiting times and delayed diagnoses.

II.LITERATURE REVIEW

Esteva et al. (2017) demonstrated in a landmark study that a deep convolutional neural network trained on 129,450 clinical images could classify skin cancer at a level of competence comparable to board-certified dermatologists. This publication established deep learning as a credible and powerful approach to dermatological diagnosis and inspired a wave of subsequent research. The study used transfer learning from Google's Inception v3 architecture, highlighting the effectiveness of leveraging pre-trained models for medical imaging tasks.

The ISIC Challenge has served as a benchmark for skin lesion classification algorithms since 2016. Codella et al. (2018) summarized results from the ISIC 2018 challenge, where top-performing systems used ensemble approaches combining multiple CNN architectures including ResNet, DenseNet, and SENet. Transformer-based architectures have also been explored, with Vision Transformers demonstrating competitive performance on skin lesion datasets. However, for deployment on resource-constrained systems, efficient architectures such as EfficientNet remain preferred due to their balance of accuracy and computational efficiency.

Transfer learning has become the dominant paradigm in medical image classification due to the limited availability of labeled medical data. Models pre-trained on ImageNet provide rich feature representations that transfer effectively to medical imaging domains. Tan and Le (2019) introduced EfficientNet, a family of models scaled using a compound coefficient that achieves superior accuracy-efficiency tradeoffs compared to previous architectures. EfficientNet-B0, the smallest variant, has been widely adopted for medical imaging applications due to its strong performance with minimal computational requirements.

Previous work on multi-class skin disease classification beyond melanoma detection has demonstrated the feasibility of classifying diverse dermatological conditions simultaneously. Tschandl et al. introduced the HAM10000 dataset containing over 10,000 labeled dermatoscopic images across 7 diagnostic categories, enabling systematic evaluation of multi-class classification approaches. More recent work using DermNet NZ and combined datasets has extended classification vocabularies to 20 or more conditions

III.METHODOLOGY

Overall Working Methodology

The overall methodology follows a standard deep learning pipeline adapted for medical image classification. The process begins with dataset collection and curation from DermNet NZ and ISIC sources. Data is cleaned, balanced, and split into training (70%), validation (15%), and test (15%) sets using stratified sampling to maintain class distribution. The preprocessing pipeline is applied consistently across all splits with augmentation applied only to the training set. The EfficientNet- B0 model is trained using the two-stage transfer learning approach described in Section 5.3. Model performance is evaluated on the held-out test set using accuracy, precision, recall, F1- score, and confusion matrix metrics. The trained model weights are saved and integrated into the Flask web application for deployment.

Efficientnet-B0 Architecture

EfficientNet-B0 is a convolutional neural network architecture developed by Tan and Le at Google Brain in 2019. The architecture is based on Mobile Inverted Bottleneck Convolution (MBConv) blocks with squeeze-and-excitation optimization. EfficientNet models are scaled from a base configuration using a compound scaling coefficient that simultaneously scales network width, depth, and resolution. EfficientNet-B0 contains approximately 5.3 million parameters, making it significantly more parameter-efficient than architectures such as VGG-16 or ResNet-50 while achieving comparable or superior accuracy.

The B0 variant uses an input resolution of 224x224 pixels and produces 1280-dimensional feature vectors before the classification head. The compound scaling approach ensures that all three dimensions of the network are scaled in a balanced manner, avoiding the performance saturation observed when scaling only a single dimension. For the DermaAI system, the final 1000-class ImageNet classification head is replaced with a 23-class head preceded by a dropout layer.

Transfer Learning Approach

Transfer learning leverages knowledge gained during training on a large source dataset (ImageNet with 1.2 million images and 1000 classes) and applies it to a smaller target domain (skin disease images). The EfficientNet-B0 backbone trained on ImageNet has learned to detect edges, textures, colors, and complex visual patterns that are directly relevant to skin disease classification. Fine-tuning these features for the dermatological domain requires significantly less data and training time than training from scratch while achieving higher accuracy due to the rich pre-trained feature representations.

Skin Detection Algorithm

The skin detection algorithm uses HSV (Hue, Saturation, Value) color space analysis to validate that uploaded images contain human skin tissue. The image is converted from BGR to HSV color space, and pixel values are checked against predefined skin color ranges: hue between 0 and 20 or 160 and 180 degrees, saturation between 48 and 255, and value between 80 and 255. If the proportion of skin-colored pixels in the image falls below a threshold of 10%, the image is rejected with an informative message asking the user to upload a clearer image of the affected skin area. This validation step prevents misclassification of non-skin images and improves the reliability of the system.

IV.MODELING AND ANALYSIS

Introduction

System analysis identifies the requirements for an effective solution by examining the limitations of existing approaches and the needs of the target users. For skin disease analysis, the core requirement is an accurate, fast, and accessible classification system that provides meaningful diagnostic support without replacing professional medical judgment. This chapter analyzes existing systems, their limitations, and the design rationale for the proposed DermaAI system.

Existing System

Existing systems for skin disease identification include manual dermatological assessment by specialists, general-purpose symptom checkers, and a limited number of AI-based mobile and web applications. Manual assessment by dermatologists provides the highest diagnostic accuracy but is limited by availability and cost. Consumer-facing applications such as DermEngine and VisualDx provide AI-assisted diagnosis but are primarily targeted at healthcare professionals, require subscriptions, and are not freely accessible. Earlier AI models for public use have been limited to binary melanoma detection or small vocabularies of 5-7 conditions, lacking the comprehensive coverage needed for general use.

Limitations Of Existing System

The major limitations of existing systems include: limited classification vocabulary restricted to a small number of conditions; lack of publicly accessible, free-to-use platforms; absence of comprehensive medical information alongside predictions; inadequate transparency in confidence reporting; failure to validate input images for skin content; and poor user interface design not suited for non-specialist users. These gaps motivate the development of a comprehensive, accessible, and informative skin disease analysis system.

Proposed System

The proposed DermaAI system is a comprehensive AI-powered skin disease analysis web application that addresses the identified limitations. The system uses EfficientNet-B0 with transfer learning to classify 23 skin conditions from uploaded images. A Flask-based web application provides an intuitive interface for image upload and result display. The system validates input images using HSV skin detection, filters predictions by confidence threshold, and provides comprehensive medical information including descriptions, symptoms, and treatment recommendations alongside each prediction. All predictions are accompanied by clear medical disclaimers emphasizing the role of the tool as a diagnostic aid rather than a replacement for professional consultation.

Advantages Of The Proposed System

The proposed system offers several significant advantages over existing approaches. It provides a broad classification vocabulary of 23 conditions covering both common and serious dermatological conditions. It is freely accessible as a web application without requiring specialized hardware or software. It provides transparent confidence scores and clear visual indicators. It integrates a comprehensive medical knowledge base providing actionable information alongside predictions. It includes robust safety features including image validation and medical disclaimers. It achieves high accuracy of 95.6% while maintaining fast inference times of under 2 seconds per image.

V.RESULTS AND DISCUSSION

Testing Environment

Model training and evaluation were conducted in a Python 3.11 environment using Google Colab with GPU acceleration (NVIDIA T4 or higher) and locally using Anaconda with CUDA support where available. Web application testing was performed on Windows and Linux systems using a standard web browser. Performance metrics were evaluated on a held-out test set that was not used during training or hyperparameter tuning. Real-world testing was conducted using photographs of skin conditions taken with standard smartphone cameras to evaluate practical deployment performance.

Model Performance Metrics

Overall Test Accuracy
95.6%
Average Precision
94.8%
Average Recall
94.2%
Average F1-Score
94.5%
Average Inference Time
<2 seconds
Model Parameters
~5.3 million
Training Dataset Size
Combined DermNet NZ + ISIC 2019
Number of Classes
23 skin conditions

Classification Results

The DermaAI model demonstrated strong classification performance across all 23 skin conditions. The model achieved the highest accuracy on conditions with distinctive visual characteristics such as melanoma, psoriasis, and tinea, where unique color patterns, texture distributions, and lesion morphologies provide clear discriminative features. Conditions with visually similar presentations, such as different types of dermatitis, showed slightly lower accuracy due to inter-class similarity.

The confusion matrix analysis revealed that the most common misclassifications occurred between visually similar conditions sharing similar inflammatory presentations. The top-3 prediction strategy ensures that the correct diagnosis is captured within the displayed results even when the top-1 prediction is incorrect, improving practical utility for screening applications. The confidence threshold of 60% effectively filters out low-reliability predictions, reducing the rate of misleading results.

Web Application Results

The DermaAI web application was successfully deployed and tested with a variety of real-world skin images. The image upload interface functioned reliably across desktop and mobile browsers. The skin detection module correctly rejected non-skin images in all test cases.

Prediction results were returned within 2 seconds for all tested images on the production server. The results page displayed predictions, confidence scores, and medical information in a clear and informative format. User testing feedback indicated that the interface was intuitive and the information provided was helpful for understanding potential skin conditions, while the medical disclaimers were clearly visible and appropriately prominent.

VI.CONCLUSION

This project successfully developed DermaAI, a comprehensive AI-powered skin disease analysis system that classifies 23 skin conditions from uploaded medical images using EfficientNet-B0 deep learning architecture with transfer learning. The system is deployed as a Flask web application providing an intuitive, accessible, and informative platform for dermatological screening support. The project demonstrates the practical feasibility and significant potential of deep learning technology to improve access to dermatological screening, particularly in resource- limited settings.

The project successfully achieved all defined objectives. The EfficientNet-B0 model achieved 95.6% overall test accuracy, substantially exceeding the target accuracy threshold. The web application provides real-time predictions within 2 seconds, meeting the responsiveness requirement for practical use. The 23-class vocabulary covers the full range of common dermatological conditions from benign lesions to serious malignancies. The integrated medical knowledge base provides actionable information alongside each prediction. All safety features including image validation, confidence thresholding, and medical disclaimers are fully implemented and operational.

The high classification accuracy, broad disease coverage, fast response times, and robust safety features collectively make DermaAI a practically viable tool for skin disease screening applications.The project demonstrates how modern deep learning frameworks, transfer learning techniques, and web technologies can be combined to create accessible healthcare tools with real-world impact. The ethical framework embedded in the system's design ensures responsible deployment while the identified future enhancement directions provide a clear roadmap for continued improvement and clinical utility.

REFERENCES

1. Esteva, A., et al. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115–118.
2. Tan, M., & Le, Q. V. (2019). EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. Proceedings of ICML 2019.
3. Tschandl, P., Rosendahl, C., & Kittler, H. (2018). The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions. Scientific Data, 5, 180161.
4. Codella, N. C. F., et al. (2018). Skin Lesion Analysis Toward Melanoma Detection: ISIC 2017– 2018 Challenge Results. arXiv:1902.03368.
5. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. Proceedings of CVPR 2016.
6. Simonyan, K., & Zisserman, A. (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv:1409.1556.
7. Redmon, J., et al. (2016). You Only Look Once: Unified, Real-Time Object Detection. Proceedings of CVPR 2016.
8. Paszke, A., et al. (2019). PyTorch: An Imperative Style, High-Performance Deep Learning Library. NeurIPS 2019.
9. Flask Documentation. Pallets Projects. Available: https://flask.palletsprojects.com/�
10. PyTorch Official Documentation. Available: https://pytorch.org/docs/stable/index.html�