IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Skin Disease Detection

Using Deep Learning

[1]Mrs. Patil Maie Dharmraj, [2]Mrs.Samb Madhuri Ashok, [3.] Prof.Dr.P.S.Kolhe

[1]Sr.Lecturer, [2]HOD, [3]HOD

[1]Electronics and Telecommunication ,

[1]TPCT´s College of Engineering, Dharashiv, Maharashtra, India

Abstract*:* Skin diseases constitute one of the most prevalent categories of medical conditions worldwide, yet timely and accurate diagnosis remains constrained by the availability of specialist dermatologists. This paper presents an automated skin disease classification system that combines a ResNet50 convolutional backbone with a two-layer Bidirectional Long Short-Term Memory (BiLSTM) network to classify dermoscopic and clinical skin images across 22 disease categories. The spatial feature map produced by the ResNet50 encoder is decomposed into a sequence of spatial patches, which are passed to the BiLSTM for sequential refinement — a design inspired by visual transformer tokenization. To enhance deployment safety, a dual-criterion Out-of-Distribution (OOD) detector rejects non-dermatological images by combining softmax confidence thresholding with normalized Shannon entropy. The system is deployed as a Flask-based web application supporting live camera capture and file upload, targeting both clinical decision-support and medical education use cases. Training employs AdamW optimization, cosine annealing learning rate scheduling, class-weighted cross-entropy loss with label smoothing, and a WeightedRandomSampler to address class imbalance. This report details the dataset, architecture, training procedure, inference pipeline, and deployment strategy.

Index Terms: Skin disease classification, deep learning, convolutional neural network, ResNet50, LSTM, transfer learning, out-of-distribution detection, medical imaging, dermatology AI.

I. INTRODUCTION

Dermatological conditions affect approximately 1.9 billion people globally at any given time [1]. Diseases ranging from benign cosmetic conditions such as acne to life-threatening malignancies like melanoma share overlapping visual features, making unaided visual diagnosis challenging even for trained clinicians. Early and accurate identification of skin lesions significantly improves patient outcomes, particularly for malignant conditions such as squamous cell carcinoma and melanoma where five-year survival rates drop sharply with delayed diagnosis [2].

Advances in deep learning have enabled image classification systems to achieve dermatologist-level accuracy on specific, narrow-scope tasks (e.g., binary melanoma detection) [3]. However, real-world clinical deployment requires systems capable of discriminating among a large number of conditions, handling highly imbalanced class distributions, and reliably rejecting images that fall outside the training distribution — such as photographs of non-skin objects, food, or scenery.

The system presented in this paper addresses these challenges through four principal contributions:

1. Hybrid ResNet50–BiLSTM architecture that leverages spatial sequence modeling over CNN feature maps to capture both local texture and global spatial dependencies within skin lesion images.

2. Dual-criterion OOD rejection combining softmax confidence and normalized Shannon entropy to prevent high-confidence misclassifications on irrelevant inputs.

3. Robust training pipeline employing transfer learning, selective layer freezing, class-weighted loss, label smoothing, and data augmentation tailored for skin lesion images.

4. Full-stack web deployment providing a responsive, browser-accessible interface supporting real-time camera capture and file upload with sub-second inference latency on Apple Silicon hardware.

II. PROBLEM STATEMENT

Skin diseases affect a large portion of the global population and early diagnosis is essential for effective treatment and prevention of severe complications. However, accurate identification of skin conditions is often challenging due to the visual similarity among different diseases and the limited availability of dermatologists, especially in remote and underserved areas. Traditional diagnostic methods rely heavily on expert clinical examination, which can lead to delays, misdiagnosis, and increased healthcare costs.

With the increasing availability of medical images and advancements in deep learning, automated image-based diagnostic systems have the potential to assist healthcare professionals and patients by providing quick and reliable preliminary analysis. However, existing automated systems often face several challenges such as handling multiple skin disease categories, dealing with imbalanced datasets, and incorrectly classifying irrelevant or non-skin images.

Therefore, there is a need to develop an intelligent and reliable automated system that can accurately classify various skin diseases from images while also identifying and rejecting irrelevant inputs. The goal of this project is to design and implement a deep learning–based skin disease detection system using a hybrid architecture combining **ResNet50** and **Long Short-Term Memory** networks. The proposed system aims to classify multiple skin conditions from dermoscopic or clinical images and provide an accessible web-based platform that supports real-time image upload or camera capture for preliminary skin disease analysis.

III. BLOCK DIAGRAM OF THE SYSTEM

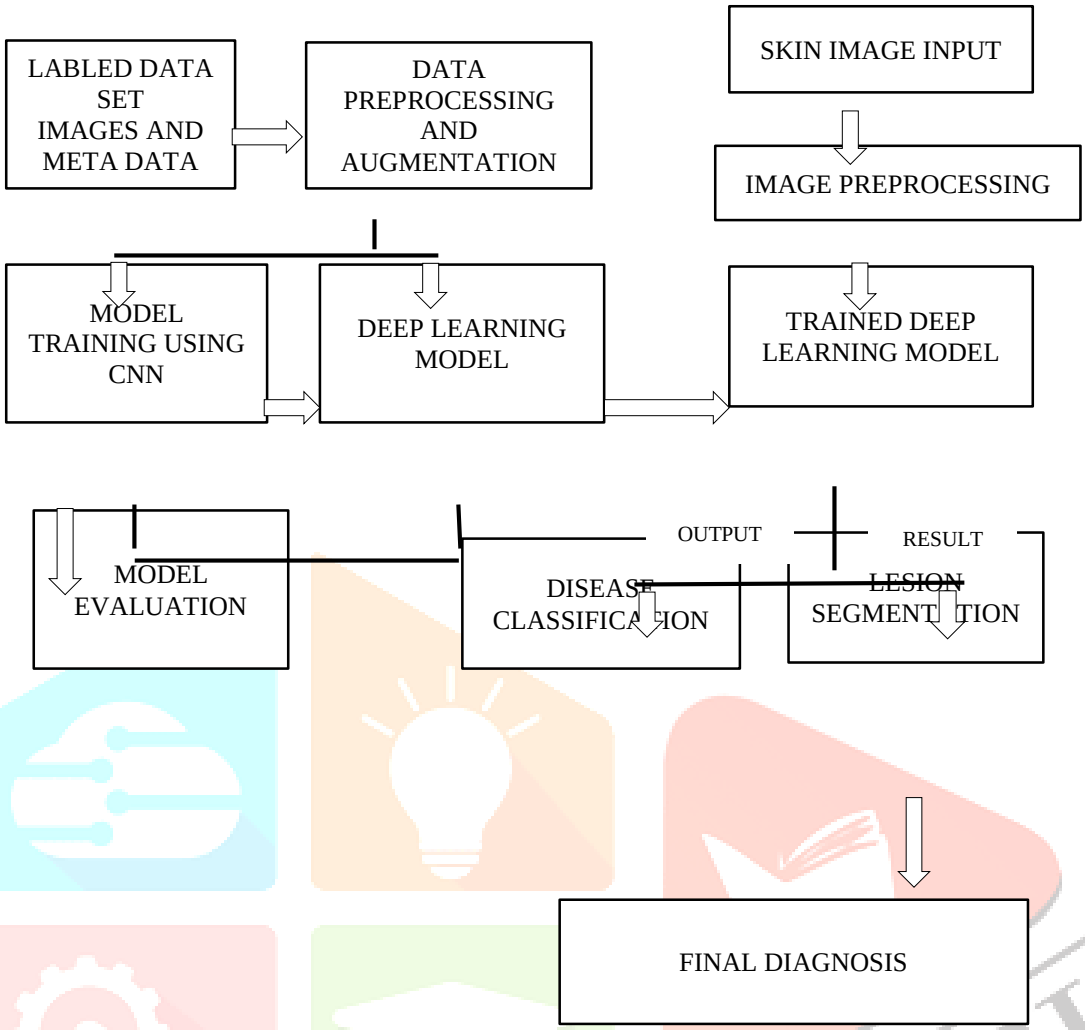

IV. RELATED WORK

A. CNN-Based Skin Lesion Classification

Early deep learning approaches to skin disease classification applied pre-trained VGGNet and AlexNet architectures fine-tuned on ISIC (International Skin Imaging Collaboration) benchmark datasets [4]. Esteva et al. [3] demonstrated that a Google Inception v3 network trained on 129,450 clinical images could classify melanoma and keratinocyte carcinoma at dermatologist parity, a landmark result that catalysed the field. ResNet architectures have since become the dominant backbone choice for medical image classification due to residual connections mitigating vanishing gradient problems in deep networks [5].

B. Recurrent Augmentation of CNNs

LSTM networks, originally designed for sequential data such as natural language, have been applied in vision tasks to model spatial context over feature sequences. Shi et al. [6] introduced ConvLSTM for spatiotemporal prediction, while several medical imaging works treat spatial rows or patches of CNN feature maps as tokens fed into an LSTM, yielding improved sensitivity for spatially complex pathologies. This spatial-patch-to-sequence paradigm predates Vision Transformers (ViT) [7] and offers a simpler implementation with comparable representational expressiveness.

C. Out-of-Distribution Detection

Hendrycks & Gimpel [8] established the baseline approach of using maximum softmax probability as a proxy for in-distribution membership, noting that in-distribution samples tend to produce higher maximum softmax outputs. Lee et al. [9] extended this with Mahalanobis distance in feature space. Our system uses a practical dual-criterion combining maximum softmax probability and normalized Shannon entropy, enabling rejection without additional calibration overhead.

D. Multi-Class Skin Disease Datasets

The ISIC Archive and HAM10000 [10] dataset are the most widely used benchmarks, covering 10 classes. The dataset used in this work extends coverage to 22 conditions including autoimmune disorders (lupus, vasculitis), fungal infections (candidiasis, tinea), vascular anomalies, and a normal/unknown class — making it substantially more clinically comprehensive.

V. CONCLUSION

This paper presented `SkinDiseaseNet`, a hybrid deep learning system for automated classification of 22 skin disease categories. The ResNet50–BiLSTM architecture extends the standard fine-tuning paradigm by treating CNN spatial feature maps as token sequences, enabling the LSTM to model inter-region spatial dependencies. A dual-criterion OOD rejection mechanism combining softmax confidence and normalized entropy provides a practical safety layer against irrelevant image submissions. The system is packaged as a full-stack web application with a clean, responsive interface supporting both camera and file-based input, making it accessible for clinical education and research without specialized software.

Future work will investigate: (1) replacing the BiLSTM head with a multi-head self-attention module for direct comparison against transformer-based architectures; (2) incorporating GradCAM visualization overlays to provide interpretable heatmaps localizing discriminative lesion regions; (3) collecting and annotating additional images for underrepresented classes (vasculitis, bullous disease); and (4) formal evaluation against HAM10000 and ISIC 2019 benchmark datasets for quantitative comparison with state-of-the-art methods.

V I. REFERENCES

[1] World Health Organization, "Skin diseases," WHO Fact Sheets, 2023. [Online]. Available: https://www.who.int

[2] American Cancer Society, "Survival rates for melanoma skin cancer," ACS, 2024.

[3] A. Esteva, B. Kuprel, R. A. Novoa, et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, 2017.

[4] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015.

[5] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE CVPR*, pp. 770–778, 2016.

[6] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-C. Woo, "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," in *Proc. NeurIPS*, 2015.

[7] A. Dosovitskiy, L. Beyer, A. Kolesnikov, et al., "An image is worth 16×16 words: Transformers for image recognition at scale," in *Proc. ICLR*, 2021.

[8] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," in *Proc. ICLR*, 2017.

[9] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," in *Proc. NeurIPS*, 2018.

[10] P. Tschandl, C. Rosendahl, and H. Kittler, "The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions," *Scientific Data*, vol. 5, p. 180161, 2018.

[11] R. Müller, S. Kornblith, and G. Hinton, "When does label smoothing help?" in *Proc. NeurIPS*, 2019.

[12] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. ICLR*, 2019.

[13] I. LOSHCHILOV AND F. HUTTER, "SGDR: STOCHASTIC GRADIENT DESCENT WITH WARM RESTARTS," IN *PROC. ICLR*, 2017