IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Deep Learning In Dermatology: End-To-End Skin Disease Detection

[1]Prajwal Vairat, [2]Omkar Patil, [3]Megha Tulse, [4]Kiran Pawar, [5]Sagar Patole

[1,2,3,4] Student's

[5] Professor

[1] Computer Engineering,

[1] PES's College of Engineering , Phaltan , Maharashtra , India

Abstract: Skin diseases are one of the most common health disorders worldwide, but early and correct diagnosis is a problem, particularly for areas with an inadequate supply of dermatologists. The current paper offers an end-to-end diagnostic system entirely based on deep learning methods for detecting and classifying different skin diseases from dermatoscopic images. With the help of convolutional neural networks (CNNs), the model learns and extracts complex skin patterns from images automatically without manual feature engineering. Having been trained on a publicly available dermatology dataset, the system shows robust performance for various skin types and disease classes. With the help of the power of deep learning, this method presents a scalable and effective solution that can help healthcare workers and possibly introduce dermatological screening systems to low-income populations. The findings discuss the revolutionary power of AI in enhancing access to correct skin disease diagnosis.

Index Terms **-** Deep Learning, Skin Disease Detection, Computer-Aided Diagnosis, Image-Based Disease Prediction, Automated Diagnosis, Medical Image Classification.

I. INTRODUCTION

Skin diseases are one of the most common health issues, impacting millions of people worldwide. Such illnesses range from mild and harmless conditions like eczema and acne to more severe and life-threatening diseases such as melanoma. Accurate and timely diagnosis is vital for successful treatment and for the prevention of complications.

Yet, access to dermatology specialists in a timely manner still poses a considerable challenge, particularly in remote and low-resource settings, tending to lead to misdiagnosis or delayed treatment. Artificial intelligence (AI) advances, especially deep learning, have lately been opening up promising integration possibilities for automating the diagnosis of skin disorders using image-based analysis. Convolutional Neural Networks (CNNs), one of the fundamental architectures of deep learning, have shown excellent performance on medical image classification, such as dermatological classification. CNNs have the capability to learn intricate patterns and features directly from images, thus lesser dependency on human intervention for feature extraction and domain knowledge. This work aims to develop a deep learning-based framework for prediction and classification of skin disease from dermatoscopic images.

We train the system on a large and varied dataset to achieve high diagnostic accuracy and wide applicability in various skin conditions and shades. The ultimate goal is to develop a low-cost, scalable, and effective tool that can assist healthcare workers in clinical decision-making and provide dermatological services to the disadvantaged.

II. LITERATURE REVIEW

The [1] paper, In 2023, Rifat Sadik, Anup Majumder, Al Amin Biswas, Bulbul Ahammad & Md. Mahfujur Rahman present this research paper to introduce a computer vision-based method that utilizes deep learning models to classify five skin diseases automatically. The researchers are able to obtain high accuracy and F1-scores by using MobileNet and Xception models with the help of transfer learning and data augmentation. A web application is created to make the system easily accessible to the users. Future works attempt to further improve the system capability with transformer-based models, expanding the dataset, and investigating ensemble methods.

The [2] paper, In 2023, Shamim Kaiser, MD, and Sazhadul Islam Protasha introduced a new approach to the detection of eczema and psoriasis based on deep learning. The authors used the robust Inception ResNet v2 model and adapted it to analyse skin images with high accuracy in the validation and testing steps. Besides developing the model, the authors also highlighted its usability. The authors created simple-to-use tools, including a smartphone app and web server, which support real-time diagnosis and severity scoring of the skin diseases. Besides enhancing the accuracy of detection of skin diseases, this work also shows how AI can be utilized to create effective and simple-to-use health care solutions.

The [3] paper, In 2024, Mr. A. Venu Gopal and his team—Achanta Sai Hari Naga Pavan, Kandula Nagendra, Mandapati Pavan Sai, and Andey Vijay Kumar—performed a research on the application of deep learning models in the detection of skin disease. They discovered that both VGG19 and the ResNetV2 base model worked well. But VGG19 surpassed them, with improved accuracy and consistent results. What made VGG19 even more remarkable was that it generalized extremely well to unseen new images and variations in skin details—a very crucial aspect for real-world application in dermatology. Their research establishes the real-world applicability of deep learning models in enabling accurate and reliable skin disease diagnosis.

The [4] paper ,In 2023, Sruthi Chintalapudi, Vikas Prateek Mishra, Shubham Sharma, and Sunil Kumar had created a predictive model for the early diagnosis of many skin diseases based on deep learning . The researchers adopted a deep learning-based methodology to improve diagnostic accuracy and broaden the model's capacity to identify a wide range of skin conditions compared to previous systems. Notably, the model was specifically designed to detect four key types of skin cancer: basal cell carcinoma, squamous cell carcinoma, benign lesions, and melanoma. What makes this approach particularly noteworthy is its ability to recognize signs of malignancy even before they become visibly apparent in patients. It was about 87% accurate in detecting early-stage skin cancers and about 91% accurate when it was tested against the training data, demonstrating its very high potential to be used for effective and early diagnosis.

The [5] paper,In 2022, Adarsh Jadhav, Shivani Hardade, Vaishnavi Phadtare, Adesh Mhetre, and Ms. Rutuja Tikait suggested a new deep learning technique based on a CNN-based model for skin disease classification. They trained and tested their model on the popularly used HAM10000 dataset, with which it attained a remarkable accuracy of 97.05%. Such high accuracy suggests that their technique can be a useful addition to clinical practice—allowing dermatologists to make earlier and more accurate diagnoses, and improving the overall quality of skin disease management and patient care.

III. PROBLEM STATEMENT

Accurate diagnosis of skin diseases remains a significant global challenge, largely due to the subjective nature of traditional diagnostic methods. In most clinical settings, diagnosis heavily depends on the visual judgment and experience of medical practitioners, which can vary widely from one clinician to another. This subjectivity often leads to inconsistencies in diagnosis, especially when distinguishing between skin conditions with similar visual characteristics—such as eczema, psoriasis, and fungal infections. The margin for error increases further in early-stage conditions, where symptoms may be subtle and easily overlooked. Another pressing issue is the global shortage of trained dermatologists, particularly in rural and low-resource regions. This imbalance leaves many patients without timely or accurate diagnoses, increasing the risk of disease progression or mismanagement.

As a result, individuals may face delayed treatment, worsening symptoms, or unnecessary anxiety—all of which impact both health outcomes and quality of life. Although advancements in medical imaging technology have enhanced the visual capture of skin conditions, most diagnostic tools still require expert interpretation. There is a critical lack of automated, scalable, and real-time diagnostic solutions that can assist healthcare providers—or empower patients directly—in identifying skin diseases accurately. One of the biggest

roadblocks in developing robust deep learning solutions lies in the limited availability of large, diverse, and well-annotated datasets. Many of the existing public datasets used to train AI models are relatively small and often lack representation across different skin tones, age groups, geographical regions, and rare skin conditions. This lack of diversity creates significant challenges in building fair and generalizable models. A model trained predominantly on images from light-skinned individuals, for instance, may underperform when diagnosing diseases on darker skin tones, leading to biased outcomes and health disparities.
In addition, real-world skin images often suffer from quality inconsistencies due to varying lighting conditions, camera types, and backgrounds.

These environmental variations can mislead AI systems unless they are specifically designed to handle such noise. Moreover, class imbalance—where common conditions dominate the dataset while rare but serious diseases are underrepresented—can skew model predictions and reduce reliability in critical cases. Furthermore, privacy concerns and regulatory constraints make it difficult to collect and share dermatological data at scale, especially from healthcare settings. This further limits the ability of researchers to build large, representative datasets needed for training high-performance AI models.

Taken together, these challenges highlight the urgent need for AI-driven skin disease detection systems that are not only accurate and explainable but also ethically trained, inclusive, and accessible. Addressing these issues is essential for developing tools that can truly democratize dermatological care and make high-quality diagnosis available to all—regardless of geography, skin type, or socioeconomic status.

IV. METHODOLOGY

To create an effective skin disease prediction system, we employed a systematic methodology combining real-world medical information and the strength of deep learning. The study began with the selection of the HAM10000 dataset, a large publicly accessible dataset of over 10,000 dermatoscopic images. The dataset includes images of seven common skin diseases, namely melanoma, nevus, and keratosis, allowing the dataset to be sufficient for training the model to distinguish among the different types of diseases. These images were preprocessed prior to training to be uniform and improve model performance.

We resized them to 224×224 pixels to meet the input size required of modern CNN architectures. We normalized the pixel values to the range 0–1 and applied data augmentation (e.g., flipping, rotation, and zooming) to mimic a wider range of clinical situations and prevent overfitting. The labels were numerically encoded to utilize with classification algorithms. For the prediction model, we used two approaches: a from-scratch custom CNN and a transfer learning model based on ResNet50, a widely used deep learning architecture pre-trained on the ImageNet dataset.In the from-scratch custom CNN, a number of convolutional and pooling layers extracted visual features from images, and then classification was done using dense layers. To generalize for the sake of avoiding the model memorizing the training data, dropout layers were incorporated. In the case of the transfer learning technique, we used the expertise of ResNet50 and fine-tuned the last layers to learn to suit skin disease classification. This made the model learn efficiently from fewer training epochs with high precision. Training was performed with the categorical cross entropy loss function, appropriate for multi-class classification, and the Adam optimizer, a fast and adaptive optimizer.

Training was performed from 30–50 epochs at batch size 32. Throughout training, we monitored key performance metrics such as accuracy, precision, recall, and F1-score to monitor how well the model was generalizing to new data. After training, we tested the model on a particular test set. A classification report and confusion matrix provided some insight into how well each class was being classified. To gain a greater insight into the model's decision, Grad-CAM visualizations were used, which show where in an image the model was looking when it was making a prediction—providing some sense of whether it was "looking" in the right places. Lastly, to close the loop between research and deployment, we deployed the trained model via an easy-to-use web interface with Flask. The demonstration system accepts skin images as input and then provides prediction outputs, demonstrating how deep learning has the potential to assist dermatologists—or even facilitate early screening in remote or underserved areas.

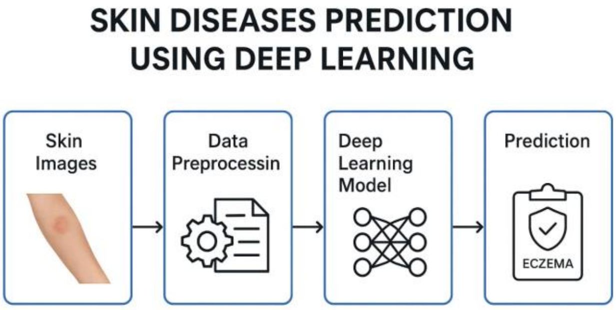

V. PROPOSED SYSTEM

In our research, we propose a new deep learning model for skin disease diagnosis that combines several state-of-the-art methods under one strong framework. At the core of our idea is a multi-task learning model, where the model can conduct three basic tasks simultaneously: lesion segmentation, disease classification, and severity estimation. This simulates the actual decision-making process of dermatologists and makes the model more clinically useful and beneficial.

Our architecture is EfficientNet B7, a high-end convolutional neural network with enhanced performance and computational efficiency. We incorporated an attention mechanism to increase the sensitivity of the system to finer, disease-specific details. This enables the model to focus on the most discriminative regions within each image—just like a trained dermatologist would examine some patterns, textures, and color differences of the skin. To make sure that the model generalizes effectively to other patient populations and imaging regimes, we employed a strong data augmentation pipeline. This includes techniques like noise injection, color jittering, and geometric transform augmentations (e.g., rotation, flipping, scaling), mimicking real-world variance and preventing overfitting. These augmentations enable the model to generalize more reliably to novel, unseen data. Training was guided by a well-designed composite loss function that integrates weighted cross-entropy, dice loss, and focal loss. Each component of the composite has a unique advantage: weighted cross-entropy handles class imbalance, dice loss improves segmentation accuracy, and focal loss focuses learning on harder-to-classify instances.
Collectively, they enable the model to learn in a stable and balanced way, especially in the context of complicated uncertain data common in dermatology. Lastly, our objective is to provide a universal AI-driven tool to assist healthcare practitioners in early diagnosis and optimal management of skin disorders. Through the provision of enhanced accuracy, generalizability, and clinical utility, we expect to minimize diagnostic delay and enhance patient outcomes, especially in communities with limited dermatological access.

VI. EXISTING WORKING / WORK

Current skin disease diagnosis models mainly rely on traditional machine learning techniques supplemented with image processing techniques to analyze the skin images. The overall system consists of acquiring high-quality dermatoscopic or clinical images, and thereafter, applying preprocessing steps to enhance image quality and homogeneity. Next, respective features such as color patterns, texture gradients, and shape descriptors corresponding to the images are manually extracted. The features are then fed as inputs to machine learning algorithms, e.g., support vector machines or random forests, for image classification into different disease classes. While these systems have been extremely promising, they also suffer from some real-world constraints. Perhaps the most significant constraint is their susceptibility to image quality—low-light images, images from other devices, or images from heterogeneous skin tones significantly impact performance.

Additionally, access to large, well-labeled medical datasets is currently a main limitation, especially those with unusual skin diseases or under-represented patient groups. The inherent difficulty of skin disease, with most conditions sharing overlapping visual features, makes classification an even larger challenge. Additionally, many existing models do not generalize reliably enough across different clinical settings, which limits their usefulness and reliability in real-world healthcare settings. In the future, there is a lot of hope that deep learning models will be able to solve such issues, and they will learn to learn useful features by themselves from raw images without any intervention from humans. Future work can even be focused on creating explainable AI models, which can make explainable decisions—something that clinicians and patients can be able to trust. Furthermore, the application of these models in mobile health apps can facilitate the easier screening for skin disease, especially in rural or underserved communities. Lastly, the exploration of multimodal approaches—unifying visual information with clinical history, symptoms, or even genomic information—can significantly enhance diagnostic accuracy and enable more personalized suggestions for treatment.

In short, while what we have now gives a good foundation, future skin disease diagnosis will be about developing more advanced, interactive, and user-friendly AI systems that bridge the gap between technology and everyday clinical practice.

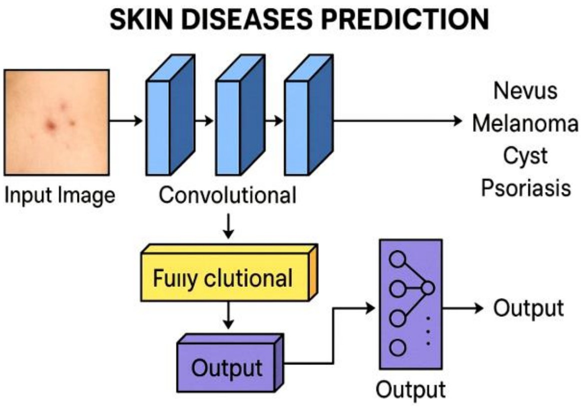

VII. RESULT AND WORKFLOW

This study presents a deep learning-based pipeline for predicting skin diseases from medical images. The process starts with collecting high-quality dermatoscopic or clinical photos of skin lesions. These images are then pre-processed through resizing, normalization, and data augmentation—techniques like rotation, flipping, and color changes are used to simulate real-world variation and boost model robustness.

After preprocessing, the images are passed through a convolutional neural network that automatically extracts relevant visual features such as texture, shape, and colour patterns. This allows the model to detect disease-specific characteristics without manual intervention.

For classification, we use powerful models like ResNet or EfficientNet, which assign a disease label and provide a confidence score for each prediction. Finally, results are displayed either visually or in a summary format, making it easy for healthcare professionals to interpret.

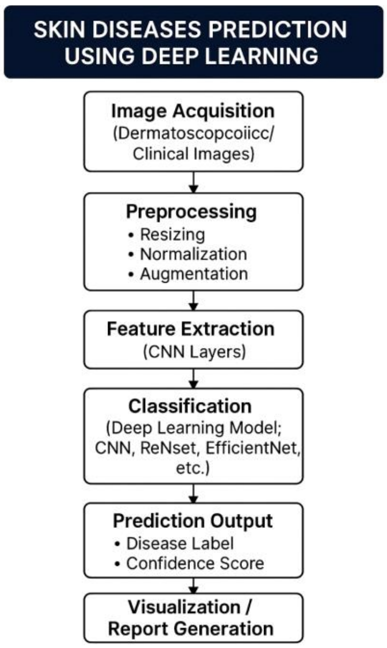

This streamlined system supports faster, more consistent diagnoses and holds promise for extending dermatological care to underserved areas.

VIII. CONCLUSION

Computer-assisted diagnosis of skin disease using deep learning can potentially greatly assist and automate clinical dermatology. In this work, we have demonstrated that the use of convolutional neural networks and specific image augmentation can greatly enhance the robustness and accuracy of skin lesion classification. Not only is this faster for diagnosis, but it presents a scalable solution that has the potential to deliver dermatological care to regions that lack specialist facilities.

Despite the system's strong performance, practical limitations remain. These include restricted access to large, diverse image datasets, the visual similarity between different skin conditions, and challenges in adapting such models to real-world medical settings. As this field progresses, future studies should aim to develop models that are more interpretable, adaptable across populations, and capable of integrating additional forms of clinical data.

By refining these methods and embedding them into healthcare infrastructure, we move closer to accessible, early-stage detection systems that can benefit both clinicians and patients—ultimately leading to improved outcomes and more efficient care delivery in dermatology.

IX. ACKNOWLEDGEMENT

Special thanks to **Prof. S.D Patole** , Assistant professor / guide for the project and **Prof. A.A Hipparkar** , Head of Department , Computer Department

X. REFERENCES

[1] An in-depth analysis of Convolutional Neural Network architectures with transfer learning for skin disease diagnosis Rifat Sadik a , Anup Majumder a , Al Amin Biswas b,∗ , Bulbul Ahammad a , Md. Mahfujur Rahman b (2023)

[2] Deep Learning-Based Skin Disease Detection Using Convolutional Neural Networks (CNN) Md. Sazzadul Islam Prottasha , Sanjan Mahjabin Farin, Md. Bulbul Ahmed, Md. Zihadur Rahman, A. B. M. Kabir Hossain, and M. Shamim Kaiser (2023)

[3] Skin Disease Detection Using Deep Learning Techniques Mr. A. Venu Gopal1*, Achanta Sai Hari Naga Pavan2 , Kandula Nagendra3 , Mandapati Pavan Sai4 and Andey Vijay Kumar5 (2024)

[4] Skin Disease Detection Using Python And Deep Learning Usha Dhankar, Seema Jain Saba Zaidi and Shubham Chauhan (2023)

[5] Skin Disease Detection using Convolutional Neural Network Srujan S A1, Chirag M Shetty2 , Mohammed Adil3 ,Sarang P K4 and Roopitha C H5 (2022)

[6] Mobile-Based Skin Disease Diagnosis System Using Convolutional Neural Networks (CNN) MWP Maduranga