



Skin Disease Detection Using CNN

Ravindra Prasad¹, Megha K², Poorvika K J³, Sandhya J V⁴, Yashaswini C K⁵

Associate Professor, Computer Science and Engineering, East West College of Engineering, Bangalore, India¹

Student, Computer Science and Engineering, East West College of Engineering, Bangalore, India²

Student, Computer Science and Engineering, East West College of Engineering, Bangalore, India³

Student, Computer Science and Engineering, East West College of Engineering, Bangalore, India⁴

Student, Computer Science and Engineering, East West College of Engineering, Bangalore, India⁵

Abstract: Skin diseases are among the most common medical conditions worldwide, affecting millions of people each year. Accurate and timely diagnosis is critical, yet traditional diagnostic methods depend heavily on dermatologists' expertise and manual examination, which can lead to human error and delayed treatment. This paper presents an automated skin disease detection system using Convolutional Neural Networks (CNNs), a deep learning technique capable of learning complex visual features from medical images. The proposed model classifies skin lesions into different disease categories, such as melanoma, eczema, and psoriasis, using publicly available datasets like HAM10000. The CNN model is trained and validated on dermoscopic images, achieving high accuracy in disease identification.

Keywords: Skin Disease Detection, Deep Learning, Convolutional Neural Network, Image Classification, Medical Diagnosis.

I. INTRODUCTION

Skin is the largest organ of the human body and acts as a protective barrier against external elements. However, it is also prone to various diseases, such as dermatitis, eczema, psoriasis, and skin cancer. Detecting skin diseases at an early stage is crucial for effective treatment and improved patient outcomes. Traditionally, diagnosis relies on clinical observation and biopsy, which are time-consuming and dependent on the experience of dermatologists. With the rapid advancement in Artificial Intelligence (AI) and Computer Vision, automated detection methods have gained significant attention. Deep learning, particularly Convolutional Neural Networks (CNNs), has revolutionized medical image analysis by automatically learning features from large datasets. CNNs have been successfully applied to detect tumors, classify X-rays, and identify retinal disorders. Similarly, in dermatology, CNNs can analyze dermoscopic images and detect disease patterns with remarkable accuracy.




Figure 1: Methodology




Figure 2: Data Flow Diagram

We have seen the growth of many technologies as medical science has advanced and many of these have enabled us to better diagnose those skin diseases. Despite this, many derms still depend on manual inspection to spot skin conditions. Though this archetypical approach shall always be the option of many practitioners, it isn't flawless. It is a time consuming, menial and prone to human error, manual diagnosis. The outcome of which is that despite the complexity of certain dermatological conditions, it is not uncommon that different specialists come to very different conclusions about a diagnosis or treatment plan. A combination of clinical experience with unpredictable skin presentations may make the continued reliance on blind manual examination a reality. Skepticism towards fully automation has also been provoked by these factors. Though the potential benefits for integration of computer based diagnostic tools into dermatology may not be so obvious, the potential benefits are worthwhile. They also available in automated systems, much faster and more accurate in diagnosis with artificial intelligence (AI) and machine learning (ML) involved. With access to huge amounts of data, these models are able to identify delicate patterns that may be passed over for the human eye, helping to lower diagnostic mistakes and variety. Rather, the goal of AI/ML systems would be to act as a powerful decision support tool rather than replacing dermatologists. They might help clinicians make more consistent and precise diagnoses, ultimately helping patients. These technologies can become novel research elements in dermatology when their components are integrated together properly. Automated diagnostic tools complement rather than replace traditional methods, bridging between human expertise and capacity to perform more routine tasks. An automation integration into the diagnostic process has the potential to transform future outcomes. In addition to the ability to revolutionize detection of skin diseases, it has the potential to improve overall patient care standards, as well as meeting healthcare delivery needs. Stepping toward these advancements better means a more speeds up and more dependable dermatological practice, one that exploits the strengths of each human in addition to computational exactness.

II. PROBLEM DEFINATION

Elements in dermatology when their components are integrated together properly. Automated diagnostic tools complement rather than replace traditional methods, bridging between human expertise and capacity to perform more routine tasks. An automation integration into the diagnostic process has the potential to transform future outcomes. In addition to the ability to revolutionize detection of skin diseases, it has the potential to improve overall patient care standards, as well as meeting healthcare delivery needs. Stepping toward these advancements better means a more speeds up and more dependable dermatological practice, one that exploits the strengths of each human in addition to computational exactness.



Skin diseases affect millions of people worldwide, ranging from mild irritations to life-threatening conditions such as melanoma. Early and accurate detection is critical for effective treatment and improved patient outcomes. Traditional diagnosis relies heavily on the expertise of dermatologists, which can be time-consuming, subjective, and prone to errors due to the variability in lesion appearance.

The objective of this study is to develop an automated system for skin disease detection using Convolutional Neural Networks (CNNs). CNNs, a class of deep learning models, have shown remarkable performance in image classification tasks due to their ability to learn hierarchical features directly from raw images. The proposed system aims to classify skin lesion images into multiple disease categories accurately and efficiently, minimizing dependency on human intervention and enabling faster diagnosis.

Skin diseases are among the most prevalent health issues worldwide, affecting people of all ages and ethnicities. They range from minor conditions, such as acne and eczema, to severe and potentially life-threatening diseases, including melanoma and other forms of skin cancer. Early and accurate detection is essential for effective treatment and improved patient outcomes. Traditional diagnosis relies primarily on visual examination by dermatologists, which can be subjective, time-consuming, and dependent on the expertise and availability of medical professionals. Misdiagnosis or delayed detection may lead to ineffective treatment and serious health consequences. Recent advances in artificial intelligence, particularly deep learning, have shown significant promise in medical image analysis.

Convolutional Neural Networks (CNNs), a class of deep learning models, are especially effective in image classification due to their ability to automatically extract hierarchical features from raw images, capturing subtle patterns and textures that may be difficult for the human eye to discern. Applying CNNs to skin lesion images can enable automated detection and classification of multiple skin disease types, reducing diagnostic errors and alleviating the burden on healthcare providers. However, developing a robust skin disease detection system presents several challenges, including high intra- class variability, where lesions of the same disease appear differently across patients, and low inter-class variability, where visually similar lesions may belong to different disease categories. Additional challenges include limited availability of annotated datasets, variations in image quality, lighting, and background, as well as the need for high clinical reliability. This research aims to design a CNN-based framework capable of accurately classifying skin lesions across multiple categories, thereby assisting dermatologists in rapid and precise diagnosis, enhancing clinical decision-making, and ultimately improving healthcare outcomes.




Figure 3: Skin diseases




Figure 4: Data Flow Diagram

III. USE CASES AND USER SCENARIOS

The integration of Convolutional Neural Networks (CNNs) in dermatology offers powerful, scalable solutions across patient care, clinical support, and public health. The primary use cases center on leveraging the CNN's image classification and feature extraction capabilities to improve accuracy and accessibility in diagnosis.

Scenario 1: Mobile-Based Early Self-Screening (Public Triage)

This scenario focuses on patient-initiated preliminary diagnosis using a smartphone application powered by an efficient CNN model (like MobileNetV2). A user, concerned about a changing mole or a new skin spot, takes a photo of the lesion. The app immediately processes the image locally or via the cloud. The CNN acts as an intelligent triage tool: if the model predicts a benign condition (e.g., a common mole) with high confidence (e.g., 98% Nevus), it advises the user to monitor the lesion and provides educational information, thereby reducing unnecessary visits to a clinic. Conversely, if the CNN identifies a high-risk pattern (e.g., 90% Malignant Melanoma), the system instantly issues a Red Alert, urging the user to seek professional consultation with a dermatologist within a specified urgent timeframe. This application dramatically improves early detection rates, especially in rural or underserved areas, by making a preliminary diagnostic tool instantly accessible.

Scenario 2: Clinical Decision Support (GP and Primary Care)

In this scenario, a General Practitioner (GP), who may lack specialized dermatological training, uses the CNN system as a diagnostic aid. When a patient presents with a suspicious lesion, the GP captures a clinical or dermoscopic image and uploads it to the hospital's secure, integrated system. The high-accuracy CNN model (such as ResNet or Inception) analyzes the image and returns a classification (e.g., "Basal Cell Carcinoma") along with an Explainable AI (XAI) feature like a heatmap (Grad-CAM). This heatmap visually highlights the specific boundaries and texture features the CNN used to reach its conclusion. The GP uses this objective AI-driven evidence to either confidently confirm a benign condition or, more importantly, prioritize and expedite a referral to a specialist. For teledermatology workflows, the same CNN system can automatically triage a queue of hundreds of patient-submitted images, pushing the highest-risk cases to the top of the dermatologist's review list, ensuring that time-critical diagnoses are addressed first.

Scenario 3: Large-Scale Public Health and Research

This use case applies CNNs to large datasets for non-individual-specific goals, focusing on research, education, and public health management. A researcher, for instance, may utilize a CNN for the rapid and automated classification of a massive dataset (e.g., 500,000 images) collected from various regional health centers. The CNN's ability to quickly and accurately classify these images into categories (Melanoma, Eczema, Psoriasis, etc.) enables large-scale epidemiological studies on disease prevalence and geographical distribution, which would be prohibitively slow with human experts alone. Additionally, in drug development and treatment monitoring, CNNs can be trained for image segmentation to quantitatively measure the area and severity of a lesion (e.g., calculating the Psoriasis Area and Severity Index - PASI) over a course of treatment, providing objective, measurable data to assess therapeutic efficacy.



Scenario 4: Automated Treatment Monitoring and Disease Surveillance

This application leverages the CNN system for objective, quantitative assessment of a patient's response to therapy and as a crucial public health surveillance tool. At the individual level, for patients with chronic skin conditions like Psoriasis or severe Eczema, the CNN is trained not only to classify the disease but also to perform semantic segmentation and severity scoring. A dermatologist or nurse captures images of the lesion at regular intervals during treatment. The CNN automatically analyzes these images, precisely calculating the lesion's area and quantifiable features like redness, scaling, and thickness, which are then used to generate an objective score (e.g., a digitized PASI score).




Figure 5: Use Case Diagram

IV. TECHNICAL IMPLEMENTATION

The technical implementation of a CNN-based skin disease detection system begins with Data Acquisition and Preprocessing, where a large, diverse dataset of labeled dermatoscopic images (like those from ISIC or HAM10000) is collected. This raw data is then rigorously preprocessed, involving resizing images to a uniform dimension, normalization to a 0-1 pixel range, and critical steps like artifact removal (to eliminate hair or bubbles) and data augmentation (rotations, flips, zooming) to expand the dataset and prevent overfitting. The core of the system is the CNN Architecture, which is typically built using Transfer Learning—fine-tuning robust, pre-trained models like ResNet or EfficientNet by replacing the final classification layers and retraining them on the skin lesion data. During Model Training, the data is split into training, validation, and test sets, and the model is compiled using optimizers like Adam and a suitable loss function (e.g., Categorical Cross-Entropy), while techniques are employed to address class imbalance.

Post-training, the model undergoes Evaluation using comprehensive metrics such as Precision, Recall, Specificity, and AUC on the unseen test set, and for clinical acceptance, Explainable AI (XAI) techniques like Grad-CAM generate visual heatmaps to justify the model's prediction. Finally, for Deployment, the model is optimized for low latency by converting it to lightweight formats like TensorFlow Lite for mobile applications, or it is deployed as a secure REST API endpoint on a cloud server to serve clinical and web-based applications, seamlessly integrating the AI into the end-user workflow.

The deployment of a Convolutional Neural Network (CNN) for skin disease detection moves beyond technical excellence into a complex landscape of clinical integration, ethical responsibilities, and regulatory compliance. A key challenge is ensuring model generalizability and equity, as the performance of a model trained on predominantly fair-skinned populations may degrade significantly for diverse skin tones, potentially introducing harmful diagnostic bias. To counteract this, models require training on highly diverse, representative datasets, and developers must address issues like class imbalance—where rare but critical malignant lesions are underrepresented—by using advanced loss functions.



like Focal Loss or employing synthetic data generation. From a clinical perspective, the system's output must be interpretable; therefore, Explainable AI (XAI) techniques like Grad-CAM are crucial to provide visual evidence (heatmaps) that justifies the AI's classification to the dermatologist, thus fostering trust and clinical adoption. Crucially, as the AI-enabled system qualifies as Software as a Medical Device (SaMD), its deployment must strictly adhere to the regulatory frameworks of bodies like the US FDA or the EU's AI Act/MDR, requiring rigorous pre-market validation, extensive technical documentation, and an ongoing, total product lifecycle approach to monitor performance, safety, and effectiveness in real-world clinical settings. This ensures not only technical accuracy but also the system's safe, ethical, and equitable contribution to patient care by acting as a reliable diagnostic assistant for both primary care physicians and specialists.

The foundational requirement is a high-quality, diverse dataset, which is often the primary bottleneck. Images must be collected to represent the full spectrum of skin phototypes (Fitzpatrick scale) to combat algorithmic bias, a critical ethical challenge where models trained primarily on fair skin underperform on darker skin, exacerbating health inequalities. Preprocessing is non-negotiable: raw images are uniformly resized, normalized, and cleaned of artifacts (like hair and rulers). The issue of class imbalance, where rare malignant lesions are severely outnumbered by common benign ones, is addressed through data-level techniques (e.g., synthetic image generation via GANs) and algorithmic solutions like Focal Loss during training, which dynamically re-weights samples to force the model to focus on misclassified and rare examples. The CNN itself is typically an advanced Transfer Learning architecture (e.g., EfficientNet or ResNet) that has its final layers retrained on the specific dermatology data, ensuring high-performance feature extraction while reducing training time.

High numerical accuracy is insufficient for clinical adoption; the system must be trustworthy. This is achieved through rigorous Evaluation, where performance is measured using clinically meaningful metrics like maximizing Sensitivity (Recall) to minimize dangerous false negatives (missing a malignancy), and demonstrating high Specificity to avoid unnecessary biopsies (false positives). Crucially, the system must address the "black box" problem via Explainable AI (XAI). Techniques such as Grad-CAM (Gradient-weighted Class Activation Mapping) generate visual heatmaps overlaid on the input image, showing the dermatologist exactly which pixels and features (e.g., asymmetry, irregular borders) drove the model's decision. This transparency is vital for a physician to validate the AI's output, integrate it into their diagnostic process, and maintain clinical accountability.

The deployment phase transitions the CNN from a research prototype to a certified medical tool. As software intended for diagnosis, it is regulated as Software as a Medical Device (SaMD) by bodies like the US FDA and the European Union (under the Medical Device Regulation, MDR, and the AI Act, which often classifies diagnostic AI as "high-risk"). This requires a comprehensive Total Product Lifecycle (TPLC) Approach, acknowledging that the model is not static. Developers must submit extensive documentation detailing everything from data provenance and bias auditing to the software's architecture and performance characteristics. For adaptive models that continuously learn in the field, regulators require a Predetermined Change Control Plan that outlines *how* the model will be safely updated and re-validated without needing a full pre-market review for every minor iteration. The core ethical commitment throughout this process is to ensure data privacy and to proactively monitor the model for drift or re-emergence of generalizability failures when exposed to new patient populations, ensuring the technology reduces, rather than amplifies, existing healthcare disparities.

The creation of a robust CNN-based skin disease detection system begins with the technical challenge of Data Acquisition and Preprocessing, requiring a large, ethnically diverse dataset to mitigate algorithmic bias (particularly on darker skin tones) and extensive data augmentation coupled with techniques like Focal Loss to counteract the severe class imbalance inherent in medical data. The model itself leverages Transfer Learning, fine-tuning powerful architectures (e.g., EfficientNet) to classify dermatoscopic images, but its clinical utility hinges on Explainable AI (XAI), where methods like Grad-CAM generate visual heatmaps to transparently justify the prediction, thereby fostering physician trust and accountability. Final deployment is governed by stringent regulatory compliance as Software as a Medical Device (SaMD), necessitating a Total Product Lifecycle (TPLC) approach for continuous safety monitoring, and often employs privacy-preserving strategies like Federated Learning to allow the model to learn from decentralized clinical data across multiple institutions without compromising patient confidentiality.

V. LITERATURE REVIEW

Skin disease images have vast variations in the texture of the image; several authors have attempted to design an automatic detection system based on the texture features. Different combinations of features with four popular ML algorithms were considered to compare classifier



performances. Classifying images of skin diseases, different types of color and texture features were used with LDA (Linear Discriminant Analysis), SVM (Support Vector Machine), ANN (Artificial Neural Networks) and KNN (K-Nearest Neighbor). LDA, SVM showed highest classification accuracy with range of 60 to 80. [1] Out of all tested classification models, LDA was better for color features in both binary and multi-class classifications and combined feature gave better accuracy only in binary classification. Dimensionality reduction and faster execution are the biggest advantages of this classifier.

Drawbacks:

- Low classification accuracy due to large feature dimension
- Proposed model gave lowest deviation from mean value.

An important step in the automated system of melanoma detection is the segmentation process which locates the border of skin lesion in order to separate the lesion part from background skin for further feature extraction. [2] This paper gives a study on various segmentation techniques that can be applied for melanoma detection using image processing. Statistical region merging, iterative stochastic region merging, adaptive thresholding, color enhancement and iterative segmentation, multilevel thresholding are discussed in this paper. A comparative study of these segmentation methods is also performed based on the parameters accuracy, sensitivity and specificity. Multilevel thresholding has the highest accuracy and specificity and maximum sensitivity is obtained for iterative stochastic region merging.

An ABCD technique is proposed to detect the malignant melanoma at an early stage in order to reduce the medical cost of taking biopsy. [4] First the skin image is filtered by using 3 Wiener filter and then segmented to extract the features by using Otsu thresholding and boundary tracing algorithm. The advantage of these methods for image segmentation is to obtain an accurate result for the feature extraction. The histogram of gradient method is used to extract the features of segmented image and then ABCD technique is applied to differentiate mole and melanoma and also find the spreading chances of melanoma.

Multi-type skin diseases classification using OP-DNN: The presented strategies focus on recognizing singular skin diseases, which formulates them hard to concern with the exact identification of multi-type skins. Something else, the current methods accomplish a lesser prediction rate. [14] Therefore, in this manuscript, a method based on OP-DNN (Optimized – Deep Neural Network) is to identify four various types of skin diseases. This classification algorithm classifies incoming clinical images as different skin diseases with the help of probability values. While learning OP-DNN, it is essential to determine the optimal weight values for reducing the training error. For optimizing weight in OP-DNN structure, an optimization approach is implemented in this research.

Digital Dermatology Skin Disease Detection Model using Image Processing: Image processing is used to detect skin diseases in humans. [15] This paper describes the current methods employed for detecting skin diseases, proposes a digital method to detect skin diseases and states the benefits of this method. Also includes a detailed description of the transforms used to implement the proposed method. The transforms implemented are compared on accuracy parameter. Prototype provides a noninvasive method of skin disease detection where the patient provides a picture of the infected area as an input to the prototype and any further analysis is done on this input image.

Image processing is used to detect skin diseases in humans. [15] This paper describes the current methods employed for detecting skin diseases, proposes a digital method to detect skin diseases and states the benefits of this method. Also includes a detailed description of the transforms used to implement the proposed method. Thus, it can be inferred that skin diseases are more concerned for human health. And it is really important to identify it in early stages in order to avoid malicious skin cancers. To overcome this problem, we developed a system that detects the skin disease based on user's input.

VI. EVALUATION AND RESULTS


The proposed CNN achieved high accuracy (92–95%) across evaluation metrics. Confusion matrices confirmed correct classification for most diseases, though minor misclassifications occurred between similar lesions. The model outperformed traditional methods such as SVM and KNN due to CNN's ability to learn hierarchical features automatically. A GUI was developed to demonstrate real-world usability.

The performance evaluation of a CNN-based skin disease detection system is critical for establishing its clinical credibility, relying on a robust set of metrics and rigorous benchmarking against human expertise. Key quantitative metrics—such as Sensitivity (Recall), which minimizes dangerous False Negatives (missed malignancies), and



Specificity, which prevents unnecessary biopsies (False Positives)—are prioritized over raw Accuracy, particularly given the often severe class imbalance in dermatological datasets. Modern models, primarily built upon advanced architectures like EfficientNet and fine-tuned using Transfer Learning, are benchmarked against public data challenges like HAM10000 and ISIC, where they consistently demonstrate high diagnostic competence, often achieving AUC-ROC scores exceeding 0.95. The most impactful result, however, is the clinical utility of the human-machine collaboration, as prospective studies confirm that a dermatologist supported by a validated CNN significantly improves their diagnostic accuracy and reduces error rates. Furthermore, the inclusion of Explainable AI (XAI) heatmaps (e.g., Grad-CAM) validates the model's decision-making process by highlighting relevant image regions, transforming the system from a black-box predictor into a trustworthy clinical decision support tool that is compliant with Software as a Medical Device (SaMD) regulations and poised for real-world deployment. In terms of architectural advancement, the field is rapidly progressing beyond traditional CNNs toward hybrid models and Vision Transformers (ViTs). Unlike CNNs, which primarily capture local features, ViTs utilize a self-attention mechanism to model long-range dependencies across the entire image, better mimicking a dermatologist's holistic, multi-scale observation of a lesion's context and global patterns. Finally, future directions emphasize multi-modal AI, where diagnostic output is strengthened by integrating image data with structured clinical information, such as patient history, age, and lesion site, thereby transitioning the system from a mere visual classifier to a truly intelligent clinical decision support tool ready for integration into teledermatology and primary care settings.

Figure 6: Sequence Diagram



VII. CONCLUSION

This study presents a CNN-based system for automated skin disease detection, achieving high accuracy and demonstrating the potential of AI-assisted dermatological diagnosis. Future work will focus on integrating multimodal data, deploying on edge devices, and enhancing performance with larger datasets.

The proposed skin disease detection method classifies the type of skin disease using complex techniques such as Convolutional Neural Network (CNN), classify the image based on the algorithm of softmax classifier and obtain the diagnosis report as an output by extracting features from raw image data of patients skin. While using an imbalanced dataset and default input data preprocessing, the accuracy was 74.73 percent. This model was deployed on the built web application and provided a 92.29 percent accuracy by employing data augmentation on preparing the input data. In future work ,the effect of data balancing for skin disease classification will be explored. Also the development of a smartphone based expert system for the multi-class skin disease classification to make the intelligent expert system accessible for people living in remote areas and with limited resources can be included.



REFERENCES

- [1]. Brinker, Titus Josef. (2018). "Skin Cancer Classification Using Convolutional Neural Networks." ,JMIR Publications.
- [2]. Malik, S.G. et al. (2024). "*High-Precision Skin Disease Diagnosis through Deep ...*"
- [3]. B. Arivuselvam (2021)" Skin Cancer Detection And Classification Using Svm Classifier".
- [4]. ALenezi, Nawal SolimanALKolifi. (2019). "A Method Of Skin Disease Detection Using Image Processing And Machine Learning." , ScienceDirect , pp. 85-92.
- [5]. Khan, M.A., et al. (2021). "*Skin lesion segmentation and classification: A unified deep learning framework of CNN and transfer learning.*"