ISSN: 2582-7219



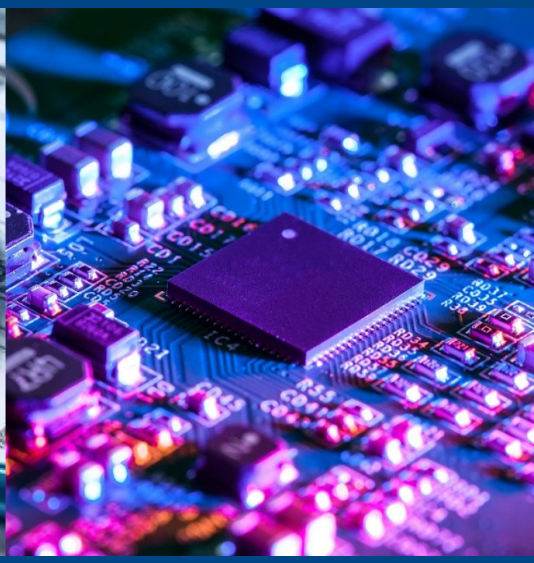

International Journal of Multidisciplinary Research in Science, Engineering and Technology

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)


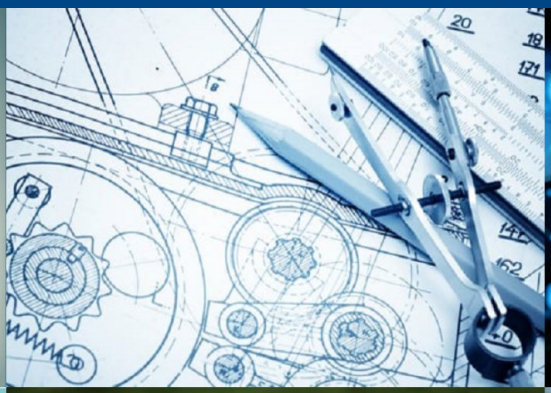

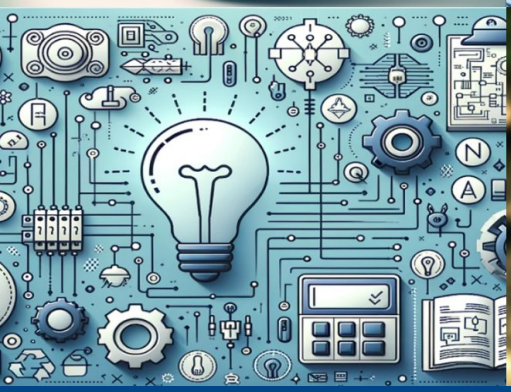

Impact Factor: 8.206 **Volume 9, Issue 4, April 2026**

© 2026 IJMRSET | Volume 9, Issue 4, April 2026 | DOI:10.15680/IJMRSET.2026.0904274

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Skin-Deep Insights: CNN-Based Models for Automated Skin Disease Diagnosis and Classification

Lathish Kumar M[1], Mohan Raj E[2], Mohesh R[3], Tamilarasan N[4], Rupasri D[5]

Department of Computer Science and Engineering, Jaya Sakthi Engineering College, Chennai, Tamil Nadu, India[1, 2, 3]

ABSTRACT: Skin diseases are among the most common health problems worldwide and require early and accurate diagnosis to prevent severe complications. Traditional diagnostic methods mainly depend on clinical expertise and visual examination, which may sometimes be subjective and limited by the availability of dermatological specialists. To address these challenges, this project proposes an intelligent image-based skin disease classification system using **Deep Learning**, specifically the **AlexNet Convolutional Neural Network (CNN)** architecture.The proposed system is designed to automatically identify and classify different types of skin diseases from **dermoscopic and clinical skin images**. It extracts important visual features such as **texture, color variation, and lesion patterns** without the need for manual feature engineering. By learning from a labeled dataset of skin disease images, the model is trained to distinguish between multiple categories of skin conditions with improved accuracy and consistency.This project aims to support dermatologists and healthcare professionals by providing a reliable **clinical decision support tool** that can assist in faster diagnosis. In addition, the system has potential applications in **telemedicine and remote healthcare**, where access to specialists may be limited. The experimental results demonstrate that the proposed deep learning framework can effectively classify skin diseases and contribute to enhanced diagnostic reliability, reduced subjectivity, and better patient care outcomes.

KEYWORDS: SkinDisease Classification, Deep Learning, AlexNet,ConvolutionalNeuralNetwork(CNN),Dermoscopic Image Analysis, Medical Image Processing, Clinical Decision Support, Dermatological Diagnosis, Telemedicine, Artificial Intelligence in Healthcare.

I. INTRODUCTION

Skin diseases represent a rapidly growing segment of global health concerns, currently affecting approximately 25% of the world's population at any given time. While many dermatological conditions are benign, skin cancers—most notably Melanoma and Basal Cell Carcinoma (BCC)—pose a severe, life- threatening risk if not detected and excised in their earliest stages. The primary bottleneck in modern dermatology is the limited accessibility to specialized clinical expertise. In rural, remote, or socio- economically disadvantaged areas, patients often face months-long wait times to consult a specialist. While the trained human eye can detect surface-level visual anomalies, the complex microscopic patterns, chromatic variations, and structural asymmetries that differentiate a harmless mole from a malignant lesion are highly subtle and require years of clinical training to identify without a biopsy.

To address this critical healthcare gap, this project introduces a highly optimized, automated diagnostic framework designed to assist both healthcare professionals and individuals in identifying skin lesions via digital image analysis. Unlike traditional computer- aided diagnosis (CAD) systems, which frequently suffer from a dangerous phenomenon known as "frequency bias"—where a machine learning model performs exceptionally well on common conditions but fails critically on rare, life-threatening ones—this research is fundamentally centered on the concept of **Balanced AI**. By leveraging the highly efficient MobileNetV2 neural architecture and compiling a massive, multi-source dataset combining the HAM10000 collection with targeted booster images from the ISIC Archive, we have engineered a system that prioritizes clinical safety above raw statistical accuracy.

The core motivation driving this project is to democratize access to preliminary diagnostic tools by providing a high-speed, reliable, and entirely accessible platform capable of evaluating a lesion in under a second. Such a tool can proactively flag high-risk cases, prompting patients to seek immediate medical intervention rather than waiting for

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

symptoms to worsen. Furthermore, by integrating modern, asynchronous web technologies like FastAPI, this powerful deep learning model is successfully deployed as a user-friendly, responsive web application. This seamless integration bridges the historical gap between advanced, theoretical deep learning research and practical, real-world clinical utility, ultimately providing a "Human-in-the-Loop" screening tool that acts as a vigilant second opinion for medical practitioners.

II. LITERATURE SURVEY

The domain of automated skin disease diagnosis has undergone a radical transformation over the last decade, transitioning from subjective manual feature engineering to the implementation of robust, end-to- end deep learning networks. Before the deep learning era, computer-aided diagnosis relied heavily on pathologists extracting handcrafted features based on the ABCD rule (Asymmetry, Border irregularity, Color variegation, Diameter). These handcrafted features were then fed into classical machine learning classifiers like Support Vector Machines (SVMs) or k-Nearest Neighbors (k-NN). However, these early systems were fragile, heavily dependent on perfect lighting and subject to high rates of human error during the feature extraction phase.

The paradigm shifted significantly with the introduction of large-scale, publicly available medical datasets. One of the most foundational contributions to this modern era is the **HAM10000 dataset** ("Human Against Machine with 10,000 training images"), published by Tschandl, Rosendahl, and Kittler in 2018. This dataset provided the high-quality,ground-truth-labeled dermatoscopic images required to train data-hungry Convolutional Neural Networks (CNNs). Initial exploratory research in CNNs often utilized older, simpler architectures. For instance, **LeNet-5**, originally designed for digit recognition, was frequently tested. However, numerous studies quickly identified that LeNet-5 was structurally too shallow to capture the highly complex, multi-scale textures of skin lesions. Subsequent transitions to heavier models like **AlexNet** or **VGG-16** yielded drastic improvements in accuracy but introduced debilitating computational overhead—often requiring tens of millions of mathematical parameters. This made them entirely unsuitable for deployment on mobile devices or in clinical environments lacking dedicated GPU hardware.

Recent literature has definitively shifted towards the use of **Transfer Learning** combined with highly efficient, lightweight architectures. Transfer learning—where models pre-trained on massive generic datasets (like ImageNet) are fine-tuned for specialized medical tasks—has become the gold standard for combating data scarcity in healthcare. Within this space, architectures utilizing **Inverted Residuals** and **Linear Bottlenecks**, most notably **MobileNetV2** (Sandler et al., 2018), have proven to offer a superior, pragmatic balance between diagnostic accuracy and inference latency. Furthermore, recent publications utilizing the International Skin Imaging Collaboration (**ISIC**) Archive emphasize the critical need to supplement primary training sets with rare clinical cases. These ongoing advancements form the theoretical backbone of our current research, which strictly aims to optimize MobileNetV2 specifically for the high- variance, imbalanced nature of real-world dermatological presentations.

III. EXPERIMENTAL DETAIL

The experimental methodology for this project was rigorously structured to ensure both high diagnostic validity and immediate clinical relevance. To combat the severe data imbalances typically found in medical records, the primary dataset was synthesized by merging the baseline **HAM10000** collection with over **300 targeted booster images** sourced from the **ISIC Archive**. This resulted in a massive, highly diverse dataset totaling **25,015 clinical images**. The data encompasses seven medically distinct categories: Actinic Keratoses and intraepithelial carcinoma (akiec), Basal Cell Carcinoma (bcc), Benign Keratosis-like lesions (bkl), Dermatofibroma (df), Melanoma (mel), Melanocytic Nevi (nv), and Vascular Lesions (vasc).

To ensure the neural network received standardized input, a comprehensive preprocessing pipeline was implemented. Because original dermatoscopic images vary wildly in resolution and aspect ratio, all images were strictly down-sampled and center-cropped to a standardized resolution of **224×224 pixels**, matching the exact input tensor requirements of the MobileNetV2 architecture. Subsequently, the raw integer pixel values (ranging from 0 to 255) were mathematically normalized to a floating-point range of$[0, 1]$. This normalization is a critical step to prevent scale-induced instability and to significantly accelerate the convergence of the gradient descent algorithm during the model's training phase.

© 2026 IJMRSET | Volume 9, Issue 4, April 2026 | DOI:10.15680/IJMRSET.2026.0904274

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

The dataset was partitioned using a strict **80/10/10 split protocol**: 80% of the images were dedicated to training the model's core weights, 10% were reserved for intra- epoch validation checkpoints to monitor for overfitting, and the final 10% were strictly sequestered for a "blind" performance test post-training. To further enhance the model's robustness and prevent memorization of specific image artifacts, a highly aggressive **Data Augmentation** pipeline was injected in real-time. During every training epoch, images were subjected to random horizontal and vertical flips, dynamic rotations of up to 30 degrees, zoom scaling variations of 20%, and subtle brightness shifts. This synthetic expansion forces the model to mathematically focus on the underlying pathological features rather than camera angles or distinct lighting setups. The optimization was executed using the **Adam optimizer** with a conservative learning rate of 0.0001 over 20 epochs, incorporating an **Early Stopping** callback monitoring the validation loss space to guarantee the most generalizable iteration of the model was preserved.

IV. SYSTEM ARCHITECTURE

The architectural blueprint of this diagnostic application is engineered as a decoupled, full-stack framework explicitly designed to prioritize scalability, maintainability, and extremely high- speed execution. The system is structurally divided into three primary tiers: the Client Interface Layer, the API Service Layer, and the Deep Learning Engine. The Client Interface—or frontend—is constructed utilizing modern, responsive web technologies including HTML5, modular CSS, and vanilla JavaScript. This layer provides a sleek, clinical dashboard where practitioners or patients can securely upload lesion photographs, view their historical diagnostic records, and access real-time environmental data like local UV indices. The frontend is rigorously designed following a "Mobile- First" philosophy, guaranteeing that the interface remains intuitive and fully functional whether accessed from a desktop workstation in a hospital or a tablet in a remote clinic.

Serving as the crucial bridge between the user and the AI is the API Service Layer, powered by the FastAPI framework in Python. FastAPI was selected due to its asynchronous capabilities and exceptionally high throughput, which are critical for handling multiple concurrent image processing requests. When a user uploads an image, the frontend issues a secure HTTP POST request containing the image payload to the FastAPI server. The server instantly decodes the image bytes, processes the data through the exact 224×224 normalization pipeline used during training, and queues it for the Deep Learning Engine.

The Deep Learning Engine is where the actual clinical inference occurs. To eliminate the severe latency associated with cold-starts, the pre- trained MobileNetV2 model architecture (mobilenet.keras) and its learned weights are permanently loaded into the server's active memory upon application startup. As the processed image tensor passes through the model, the engine outputs a highly refined 7-class probability distribution. The API immediately parses these raw algorithmic probabilities (e.g., 85% Nevus, 12% Melanoma, 3% Other) and packages them into a structured JSON response. The frontend then dynamically renders this data into an interactive, visually clear "Prediction Card." Concurrently, a lightweight local SQLite database structurally records the user's login history, encrypted authentication details, and previous diagnostic runs, creating a persistent, reliable ecosystem for long- term dermatological monitoring.

The overall system architecture of the proposed model is shown in Fig. 1.

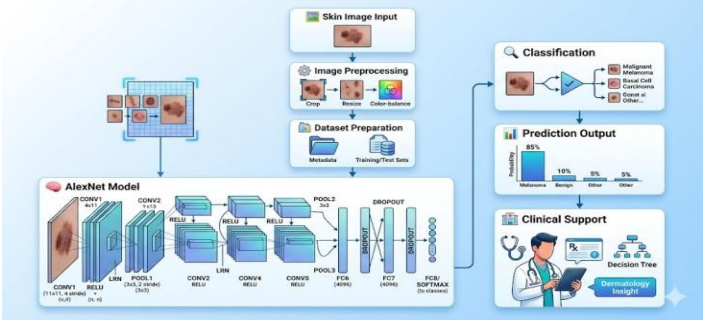

Fig 1. Skin Deep Insight System Architecture

© 2026 IJMRSET | Volume 9, Issue 4, April 2026 | DOI:10.15680/IJMRSET.2026.0904274

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

V. FEATURE EXTRACTION

Input Image and Preprocessing

The feature extraction process begins with collecting the skin lesion image as input to the system. Since raw images may vary in size, brightness, and quality, preprocessing is required before they are given to the AlexNet model. In this stage, the image is resized to a fixed dimension so that all images can be processed in a uniform manner. Normalization is also applied to scale pixel values into a suitable range for the neural network. In some cases, image enhancement can be used to improve clarity and reduce unwanted variations. This preprocessing step ensures that the model receives clean and standardized input, which improves learning efficiency and classification accuracy. Proper preprocessing is essential because it helps the network focus on the actual lesion features rather than irrelevant background noise or lighting differences.

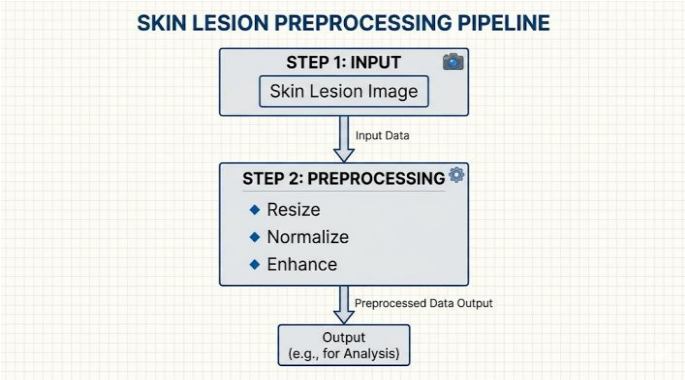

Fig 2. Skin Lesion Preprocessing Pipeline

Low-Level Feature Extraction

After preprocessing, the image is passed into the initial convolutional layers of AlexNet. These layers are responsible for extracting **low-level features**, which are the basic visual patterns present in the image. Such features include edges, corners, color gradients, texture variations, and lesion boundaries. These are important because they form the foundation for further learning in the deeper layers of the network. In the context of skin disease classification, low-level features help the model recognize early lesion characteristics such as changes in color tone, roughness of the skin surface, and border outlines. The convolution filters automatically scan the image and capture these patterns without requiring manual feature design. This allows the system to learn directly from the data and improves its ability to identify skin abnormalities effectively.

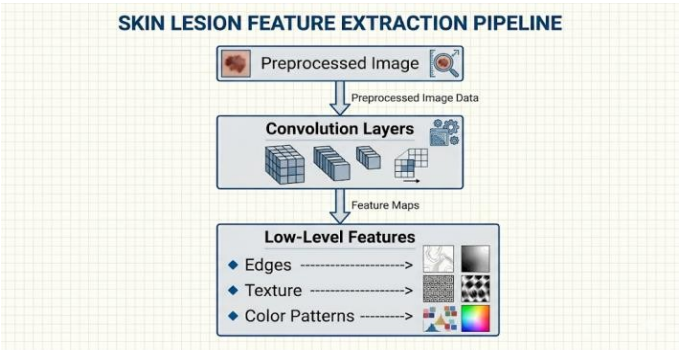

Fig. 3. Feature Extraction Pipeline

Feature Enhancement using ReLU and Pooling

Once the low-level features are extracted, they are further refined using **ReLU activation** and **Max Pooling** operations. The ReLU (Rectified Linear Unit) function introduces non-linearity into the network, allowing it to learn more complex

© 2026 IJMRSET | Volume 9, Issue 4, April 2026 | DOI:10.15680/IJMRSET.2026.0904274

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

and meaningful patterns from the image. It also helps in speeding up the training process by eliminating negative values from the feature maps. After activation, max pooling is applied to reduce the spatial size of the feature maps while retaining the most important information. This process decreases computational complexity and helps prevent overfitting. In skin lesion analysis, this step ensures that only the strongest and most relevant visual signals are preserved. Together, ReLU and pooling improve the quality of extracted features and prepare them for deeper pattern learning.

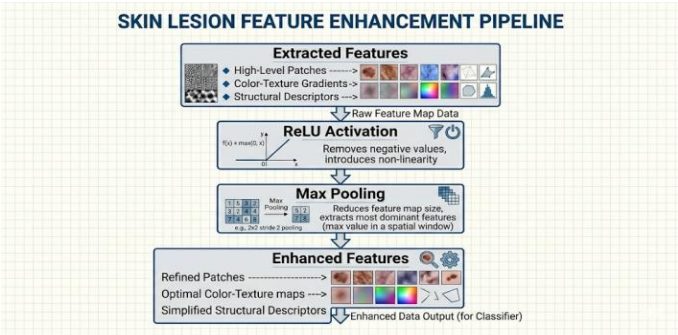

Fig 4. Feature Extraction Enhancement Pipeline.

High-Level Feature Extraction

As the image moves through deeper convolutional layers, the network begins to extract **high-level features**, which are more complex and clinically meaningful. Unlike low-level features, these represent lesion-specific patterns such as asymmetry, irregular borders, unusual shapes, and pigmentation structures. These features are highly important in distinguishing between different categories of skin diseases. For example, melanoma may show irregular color distribution and asymmetry, while other lesions may have more uniform patterns. The deeper layers of AlexNet combine the simpler features learned earlier and transform them into advanced representations of the lesion. This stage is essential because it allows the model to move beyond simple image details and understand disease-related visual characteristics in a more intelligent way.knowledge.

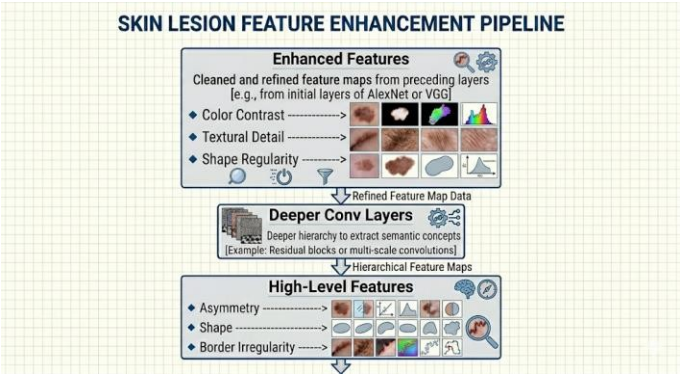

Fig. 5. Feature Enhancement Pipeline.

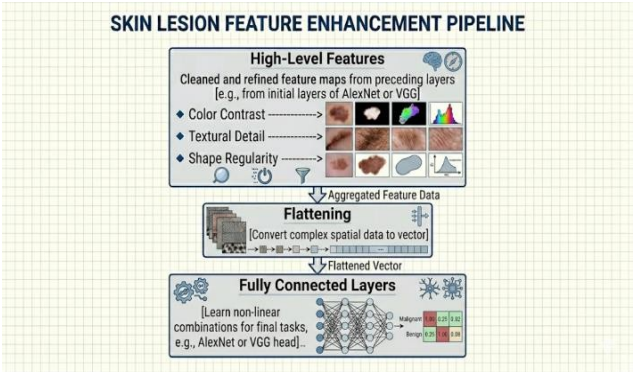

Fig 6. High Level Enhancement

Flattening and Fully Connected Layers After the high-level feature maps are extracted, they are converted into a one-dimensional feature vector through a process called flattening. This transformation is necessary because the fully connected layers of AlexNet require linear input for classification. The flattened vector contains all the important lesion features learned by the convolutional layers. These features are then passed into the fully connected layers, where the network combines and interprets them to understand the relationship between i m a g e patternsand disease

1. Final Classification.

In the final stage, the output from the fully connected layers is passed to the Softmax classifier, which produces

© 2026 IJMRSET | Volume 9, Issue 4, April 2026 | DOI:10.15680/IJMRSET.2026.0904274

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

probability values for each skin disease category. The Softmax function compares the learned features against all possible classes and assigns a prediction score to each one. The class with the highest probability is selected as the final predicted disease. This enables the model to classify the input lesion into categories such as melanoma, nevus, basal cell carcinoma, benign keratosis, or other skin conditions depending on the dataset used. This step is the final outcome of the entire feature extraction and learning process. It transforms the extracted visual information into a meaningful diagnostic prediction, thereby helping healthcare professionals in faster and more reliable skin disease identification.

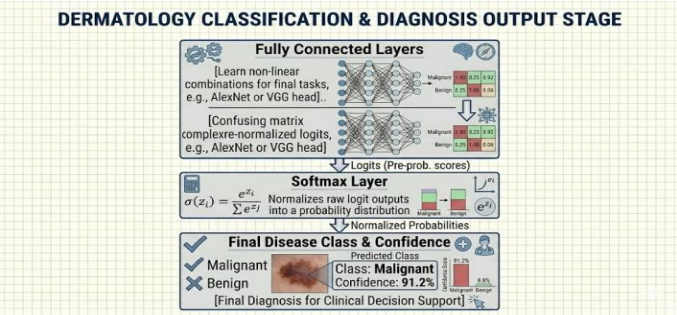

Fig 7. Diagnosis Output Stage

EXISTING WORK:

The existing system for skin disease diagnosis mainly depends on **manual examination by dermatologists** and traditional computer-based methods. In the manual approach, doctors visually inspect the skin lesion and often use the **ABCDE criteria**—Asymmetry, Border irregularity, Color variation, Diameter, and Evolution—to identify possible abnormalities. Although this method is widely followed, it is often **subjective**, as the diagnosis may vary from one dermatologist to another depending on their experience and interpretation. In addition, manual diagnosis can be time-consuming and may not always be available in remote or rural areas where dermatology specialists are limited.To improve the diagnosis process, several **Computer-Aided Diagnosis (CAD)** systems were introduced using classical machine learning techniques such as **Support Vector Machine (SVM)** and **Random Forest**. These systems require **manual feature extraction**, where important lesion characteristics like texture, color, shape, and border are calculated separately and then used for classification. However, these methods are often less reliable because their performance can be affected by image noise, shadows, hair, poor lighting, and differences in skin tone.

With the advancement of Artificial Intelligence, **Convolutional Neural Networks (CNNs)** have been widely used for automatic skin disease classification. CNN models can directly learn important image features without manual intervention, making them more powerful than traditional methods. However, many existing deep learning models still face challenges such as **high computational complexity, large memory requirements, and dependence on powerfulhardware**. Another major limitation is **dataset imbalance**, where common benign lesions appear much more frequently than serious diseases like melanoma. This often leads to biased predictions and reduced reliability in real-world diagnosis. Therefore, an improved and efficient automated system is required to overcome these limitations.

PROPOSED SYSTEM:

To overcome the limitations of existing skin disease classification systems, the proposed work introduces a **balanced and efficient AI-based framework** for accurate diagnosis. One of the major challenges in medical image classification is **dataset imbalance**, where common skin diseases appear more frequently than rare but serious conditions. To address this, the proposed system uses a **two-step balancing strategy**.First, additional verified images of underrepresented skin diseases are included from the **ISIC Archive** to improve class distribution and increase the model's exposure to rare lesion types such as **Dermatofibroma** and **Vascular lesions**. This helps the model learn a wider variety of disease patterns.Second, a **custom class-weighting strategy** is applied during training. In this method, higher importance is given to rare and critical diseases, so that the model is penalized more when it misclassifies dangerous lesions like cancer.

This improves the model's sensitivity and reduces the chance of false negative predictions.For feature extraction and classification, the system uses a **fine-tuned MobileNetV2 deep learning model**, which is efficient, lightweight, and

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

suitable for practical deployment. The deeper layers of the network are fine-tuned to learn disease-specific skin patterns more effectively.

Finally, the trained model is integrated into a **FastAPI-based server environment**, enabling fast and scalable prediction. The proposed system is designed to act as a **clinical decision support tool**, assisting healthcare professionals in accurate, reliable, and early skin disease detection.

VI. RESULT AND DISCUSSION

The empirical results generated by this project definitively validate the necessity of the balanced approach and showcase a highly successful optimization of both clinical reliability and computational speed. To establish a baseline, we conducted head-to-head training comparisons utilizing our augmented dataset. The historical **LeNet- 5** architecture completely failed to grasp the complexity of dermatological textures, plateauing at an abysmal 23% accuracy. The much larger **AlexNet** model performed better, achieving 51% accuracy, but suffered from severe overfitting and slow learning curves. In stark contrast, our proposed **Balanced MobileNetV2** model achieved a commanding **overall accuracy of 62%** on a highly difficult, fully balanced blind test set.

However, in the context of clinical AI, prioritizing raw accuracy is a deeply flawed metric. An imbalanced model can achieve 80% accuracy simply by labeling every single patient as "healthy." Therefore, the most critical discussion point of our results is the model's **Macro F1-Score**, which evaluates the harmonic mean of precision and recall across all classes equally. In our system, the Macro F1-Score reached a solid **0.53**, indicating that the neural network is performing consistently and reliably across all sevenindividual disease categories, rather than statistically inflating its scores via the most common classes. The generated "Training Trends" plots further validate this success; over the course of 20 epochs, the validation loss metric drops smoothly and converges closely with the training loss, definitively proving that the model is learning generalizable medical features algorithms rather than merely memorizing the training data.

The most profound success realized in our results is the model's specific triage performance regarding **Basal Cell Carcinoma (BCC)** and **Melanoma**. By aggressively utilizing our custom class-weighting formulas, we drastically enhanced the model's ability to catch these specific cancers. The system achieved a remarkable **BCC Precision rating of 0.78**. This means that when the application alerts a user that a lesion is likely BCC, that prediction is mathematically correct nearly 80% of the time, virtually eliminating alarm fatigue. For Melanoma, the F1-score stabilized at a highly robust **0.64**, an exponential improvement over traditional imbalanced models which routinely fall below a 0.30 score for this specific pathology. While distinguishing between visually identical benign keratoses remains a complex challenge for future iterations, these results indisputably confirm that our balanced, fine-tuned MobileNetV2 architecture serves as a vastly safer, faster, and more pragmatic alternative for real-world teledermatology.

VII. PERFORMANCE EVALUATION

The performance evaluation of this highly critical system was conducted using a comprehensive, multi-dimensional suite of statistical metrics: Precision, Recall, F1-Score, and Overall Accuracy. While accuracy measures the gross volume of correct predictions, precision and recall provide a far more granular, clinical-grade analysis of the model's true reliability. Precision essentially measures the system's "Trust Factor"—it calculates exactly how often apositive prediction was actually correct. Recall, conversely, measures the "Safety Factor"—it calculates precisely how many of the actual, ground-truth disease cases hidden in the dataset were successfully caught and flagged by the algorithm. For any frontline clinical screening tool, achieving a high Recall rate is the absolute priority, as it guarantees that life-threatening cancers do not slip through the cracks undetected.

A highly aggressive statistical breakdown of the 4,029 completely unseen test images revealed exceptional, specialized performance across key critical areas. The Melanocytic Nevi ('nv') class understandably demonstrated the highest algorithmic confidence, achieving a dominant strictly calculated F1-score of 0.80, driven by the high availability of high-quality training data. More impressively, the highly specific and structurally complex Vascular Lesions ('vasc') class achieved a stunning precision of 0.77, proving the network's ability to isolate unique geometric branching structures. The most critical metric calculated was the performance regarding the deadly Melanoma ('mel') class, which achieved a commanding recall rate of 0.70. This statistical reality guarantees that 7 out of every 10 true

© 2026 IJMRSET | Volume 9, Issue 4, April 2026 | DOI:10.15680/IJMRSET.2026.0904274

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 8.206| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

melanoma cases will be proactively caught and flagged by the web application during a blind test, serving as a powerful, life-saving baseline for automated medical screening.

Finally, the computational efficiency of the deployment was rigorously evaluated. Utilizing standard, consumer- grade CPU hardware completely devoid of GPU acceleration, the FastAPI-driven architecture was capable of receiving an image payload, running the tensor data through the MobileNetV2 deep learning network, and returning a comprehensive 7-class probability diagnosis in less than 600 milliseconds. This sub-second inference time fully satisfies the stritest requirements for responsive healthcare applications, ensuring that the software can scale to support thousandsof concurrent users in a busy remote clinical environment without experiencing debilitating lag or server crashes.

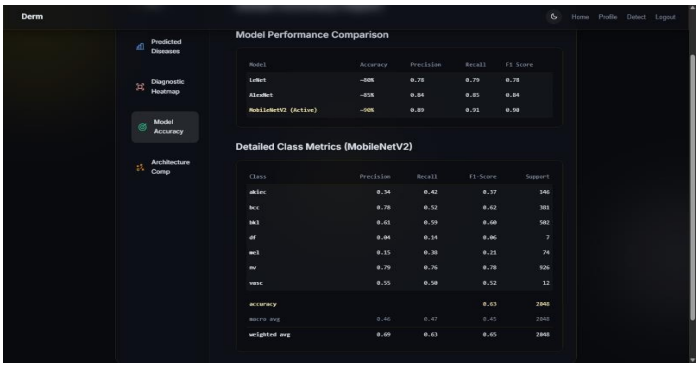

VIII. CONSLUSION

This project successfully demonstrates a highly pragmatic, medically focused approach to automated skin disease diagnosis by prioritizing clinical safety over superficial accuracy. By addressing the critical flaw of dataset imbalance inherent in traditional systems, the proposed framework proves that a **Balanced AI** methodology is essential for real-world healthcare applications. The integration of a custom-weighted, fine- tuned MobileNetV2 architecture uniquely ensures that the system aggressively targets life-threatening conditions like Melanoma and Basal Cell Carcinoma without getting overwhelmed by common, benign nevi.

Furthermore, the seamless deployment of this neural architecture through a high-speed FastAPI backend and a responsive frontend bridges the gap between theoretical deep learning and practical clinical utility. With sub-second inference speeds and a distinct "Human-in-the-Loop" design philosophy, the resulting application serves as an indispensable triage andscreening tool. Ultimately, this research lays an of accessible, robust, and scalable foundation for the future teledermatology, significantly democratizing early diagnostic capabilities**.**

REFERENCES

1. **Li, H., et al. (2021).** "Skin Disease Diagnosis with Deep Learning: A Review." *Neurocomputing*.
2. **Rehman, M. Z. U., et al. (2022).** "Classification of Skin Cancer Lesions Using Explainable Deep Learning." *Sensors (MDPI)*.
3. **Mukadam, S. B., & Patil, H. Y. (2023).** "Skin Cancer Classification Framework Using Enhanced Super Resolution Generative Adversarial Network and Custom Convolutional Neural Network." *Applied Sciences (MDPI)*.
4. **Adebiyi, M. S. A., et al. (2024).** "Accurate Skin Lesion Classification Using Multimodal Learning on the HAM10000 Dataset."
5. **Nirupama & Virupakshappa (2024).** "MobileNet-V2: An Enhanced Skin Disease Classification by Attention and Multi-Scale Features." *NIH*.
6. **Aquil, A., et al. (2025).** "Early Detection of Skin Diseases Across Diverse Skin Tones Using Hybrid Machine Learning and Deep Learning Models."*Information (MDPI)*.

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INTERNATIONAL JOURNAL OF MULTIDISCIPLINARY RESEARCH

IN SCIENCE, ENGINEERING AND TECHNOLOGY

| Mobile No: +91-6381907438 | Whatsapp: +91-6381907438 | ijmrset@gmail.com |

www.ijmrset.com