AVE Trends in Intelligent Health Letters

Artificial Intelligence-Based Facial Skin Disease Detection and Personalized Care Assistant

G. Brintha[1], Amira Sanoli[2], M. Arline Jemi Flora[3,*], Madhan Raj Gopi Akila[4], S. Tejas[5]

[1,2,3]Department of Artificial Intelligence and Data Science, Francis Xavier Engineering College, Tirunelveli, Tamil Nadu, India.
[4]Department of Computer Science, University of Stuttgart, Stuttgart, Baden-Württemberg, Germany.
[5]Department of Data Science, Analytics and Engineering, Arizona State University, Tempe, Arizona, United States of America.
gbrinthanagarajan@gmail.com[1], amirasanoli149121@gmail.com[2], arlinejemi15@gmail.com[3],
st199709@stud.uni-stuttgart.de[4], tsundar1@asu.edu[5]

*Corresponding author

Abstract: The AI-Based Skin Disease Detection and Personalised Care Assistant aims to make skin care easier to access, more accurate, and easier to use for people in both cities and rural areas. The system uses a Flutter-based mobile app and a Python Flask backend to make it easy to handle data and analyse it in real time. The app can analyse pictures of skin problems taken directly with a smartphone camera using deep learning models such as MobileNetV2. The system uses OpenCV techniques to assess the severity of detected conditions, enabling it to distinguish between mild, moderate, and severe cases. The assistant uses this information to make personalised care suggestions, which help users learn more about their condition and decide what to do next. This helps people assess their own health, reduces unnecessary hospital visits, and makes dermatologists' jobs easier. For more serious cases, GPS integration lets the system quickly connect users with dermatology specialists in their area, ensuring they receive medical help right away. There are strong security measures in place to keep private medical information safe and protect patient privacy. Overall, the solution creates a safe, smart, and easy-to-use digital platform. It helps with early detection, supports informed decision-making, and encourages better skin health management worldwide through modern AI technologies.

Keywords: Severity Estimation; GPS Integration; Deep Learning (DL); Skin Disease Detection; Skin Health Management; Modern AI Technologies; Decision-Making; Dermatologists' Jobs.

Cite as: G. Brintha, A. Sanoli, M. A. J. Flora, M. R. G. Akila, and S. Tejas, "Artificial Intelligence-Based Facial Skin Disease Detection and Personalized Care Assistant," *AVE Trends in Intelligent Health Letters*, vol. 3, no. 1, pp. 53–67, 2026.

Journal Homepage: https://avepubs.com/user/journals/details/ATIHL

Received on: 24/02/2025, **Revised on:** 18/06/2025, **Accepted on:** 26/09/2025, **Published on:** 05/01/2026

DOI: https://doi.org/10.64091/ATIHL.2026.000265

1. Introduction

The skin is the body's first line of defence against harmful microorganisms, UV rays, and physical damage. It can get a lot of different medical problems, from common inflammatory conditions like Acne and eczema to more serious ones like melanoma, which can kill you. Skin-related illnesses make up a large part of the world's non-fatal disease burden, affecting hundreds of millions of people around the world. But skin problems are often not taken seriously and are seen as cosmetic issues that don't need immediate treatment [4]. In modern healthcare systems, standard clinical practices mostly rely on a dermatologist's visual examination, dermoscopic analysis, and, when there is doubt, invasive biopsy procedures. This method is still clinically reliable, but it is not very efficient in practice and is hard to access. One major problem is the shortage of trained dermatologists, especially in rural areas and developing countries. Limited access to specialists leads to long wait times, higher travel costs, and the need to see non-specialists, which can delay diagnoses and reduce the effectiveness of treatment. Conventional diagnostic methods are intrinsically subjective and shaped by the clinician's knowledge, experience, and contextual variables [19].

Visual similarities among various skin conditions can hinder precise identification, while human factors such as fatigue and cognitive bias may exacerbate diagnostic discrepancies. Early detection is very important for a patient's survival, and failing to detect the disease early can greatly reduce their chances. Recent advances in Deep Learning and Artificial Intelligence have opened new ways to change how medical tests are performed. Dermatology is a great field for computer vision applications because it mostly uses pictures. Convolutional Neural Networks have shown they can detect subtle visual patterns, textures, and colour changes in medical images that the human eye might miss. Empirical research indicates that proficiently trained AI models can equal or exceed the diagnostic efficacy of seasoned dermatologists in regulated settings. To address these problems, this paper proposes an AI-Based Skin Disease Detection and Personalised Care Assistant, a mobile-focused healthcare platform. The suggested system uses the MobileNetV2 architecture, a lightweight convolutional neural network that performs well on mobile devices. The system has a smart severity assessment module that uses OpenCV and image processing methods to analyse colour distributions and lesion coverage in the HSV colour space and estimate the severity and size of skin lesions. The system does more than identify diseases; it also supports all aspects of lifestyle and care management, such as personalised dietary advice and daily skin care routines tailored to the diagnosed condition [20].

2. Related Work

Over the last 10 years, research into automated recognition of skin diseases has made significant progress. This is mostly because of new developments in deep learning and the growing number of annotated medical image datasets. One of the first and most important studies showed that deep convolutional neural networks (CNNs) could classify melanoma and non-melanoma skin cancers with the same level of accuracy as a dermatologist. This demonstrated that AI could be used as a diagnostic tool rather than solely as a research tool [1]. This work sparked more research into CNN-based frameworks for a wider range of skin conditions, not just cancer detection. Now, they are also used to identify common conditions such as Acne, eczema, and psoriasis. Datasets that are open to the public, such as the ISIC archive and the HAM10000 collection, have been very helpful for this research because they provide large, well-annotated image repositories from diverse sources. This allows systematic comparison of algorithms under the same conditions [3]; [5]. These datasets helped set performance standards and prompted the community to focus on how well models perform and how well they handle changes in lighting, skin tone, and lesion shape. As classification tasks became increasingly intricate, researchers examined the comparative advantages of various network architectures. AlexNet and other modified deep networks, which are traditional CNN variants, showed good results for classifying skin diseases. However, they had many parameters and were expensive to run, making them less useful in resource-limited settings, such as mobile health apps [10]; [13].

In response, the creation of efficient architectures such as MobileNet and its improved version, MobileNetV2, introduced depthwise separable convolutions and inverted residual blocks, reducing computational load while maintaining high accuracy. This makes them great for real-time inference on handheld devices [2]; [12]. Research utilising MobileNetV2 and analogous lightweight models for dermatological image classification has validated that these architectures achieve a pragmatic equilibrium between performance and deployability, thereby endorsing the viability of mobile-centric diagnostic instruments [8]. Additional research has investigated combining traditional image processing with deep learning to improve feature extraction and pre-processing before classification. The OpenCV library and other tools have been widely used for tasks such as lesion segmentation, colour space transformation, and noise reduction. These tasks help separate areas of interest and make the model focus more on clinically important features [7]. This hybrid approach is especially useful when images are taken in less-than-perfect conditions, such as when the lighting changes or there is a lot of background noise, which is common in photos taken by patients on their smartphones. Transfer learning has also become a useful approach to handling limited labelled data in medical imaging, alongside architectural choices. Researchers have improved convergence and classification accuracy by pre-training networks on large general image datasets and then fine-tuning them on domain-specific skin lesion images. This is especially useful when there isn't enough data to train a strong model [6].

Comparative studies further demonstrate that deep learning models can equal or surpass dermatologists' performance on specific diagnostic tasks. For instance, in controlled assessments of dermoscopic melanoma detection, CNN-based systems exhibited diagnostic accuracy comparable to that of numerous practising dermatologists, highlighting the potential of AI as a scalable complement to clinical expertise [9]; [18]. In addition to basic classification, researchers have looked into attention mechanisms and other architectural improvements to help models focus on the most discriminative visual patterns in skin images. This has made them better at tasks that require more detailed classification [11]. Systematic reviews of AI in dermatology highlight both the potential and the difficulties of these methodologies, indicating that although model accuracy has markedly improved, challenges such as dataset bias, insufficient diversity, and restricted generalisability across demographic groups persist as significant issues [14]. Dataset bias can lead to performance differences when models are mostly trained on images of specific skin tones or lesion types. This has led to calls for greater openness in data collection and evaluation methods. Researchers have examined how the Internet of Things (IoT) and mobile technologies can work together with algorithmic improvements to support continuous monitoring and patient engagement beyond traditional clinical settings.

IoT-based skin-monitoring frameworks demonstrate how sensor data and regular image capture can enhance longitudinal health tracking and early-detection initiatives, broadening the scope of AI from singular image classification to continuous care support [15]. The rapid rise in the use of cross-platform mobile development tools like Flutter has made it easier to build health apps that are easy to use and can connect to AI services, leveraging built-in device features like cameras and GPS. This has made diagnostic tools more available to regular users [17]. The World Health Organisation's focus on neglected skin diseases underscores the importance of diagnostic technologies that many people can use to meet public health needs, especially in areas where specialist care may not be available or be limited [16]. Together, these studies show a clear path toward lightweight, efficient, and clinically useful systems for detecting skin diseases that leverage powerful deep learning models and are easy to use in real life. Research consistently shows that AI-based methods can classify data very well and could be used in real time on mobile devices. Still, issues with data diversity, ease of understanding, and clinical validation are being investigated. Adding severity estimation, user-centred design, and secure data handling to these systems makes them even more useful in real life. This points to a future where smart diagnostic assistants help both patients and healthcare professionals with everyday care.

3. Methodology

The research methodology outlines the development of a scalable, user-friendly AI system for detecting skin diseases and assessing their severity, leveraging deep learning and computer vision techniques. It includes steps such as gathering high-quality data, pre-processing it systematically, selecting the best model, training it effectively, estimating severity, and deploying it via a mobile app. The system captures users' skin images, processes them to improve appearance, uses a lightweight deep learning model to sort them, and leverages transfer learning to speed things up. A severity estimation module assesses infection severity, and the deployed model uses a cloud backend for real-time inference, enabling users to easily view and understand the results.

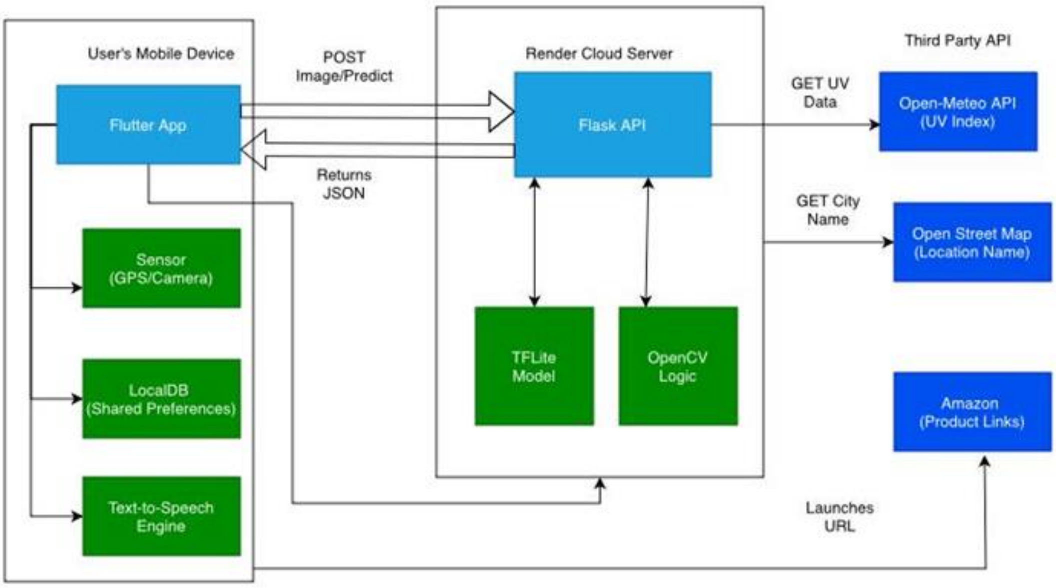

Figure 1: Block diagram of the proposed system

Figure 1 shows the overall system architecture of the SkinSutra AI app. It shows how the mobile app, cloud server, and third-party services work together to provide intelligent skin analysis and personalised care. The Flutter app is the main interface on the user's mobile device. It lets users take pictures of their skin, access GPS data, save their preferences locally, and receive spoken feedback via a text-to-speech engine. These features make the app fun to use, easy to get to, and good for everyday use. When a user sends an image, the Flutter app sends a secure POST request to the Flask API that is hosted on a cloud server. This server does the main processing and connects the mobile client to the smart analysis modules. A TensorFlow Lite model inside the server classifies skin diseases, and OpenCV logic is used to process images and determine their severity. The mobile app gets a structured JSON response with all the results. The Flask API also interacts with other services to improve the analysis. It gets location names from OpenStreetMap and UV index data from the Open Meteo API using GPS coordinates. You can also use Amazon services to open links that are related to products or care. This architecture ensures computations run quickly, feedback is delivered in real time, data is handled securely, and the user experience is smooth. It also supports reliable, scalable mobile dermatology assistance for modern healthcare delivery, improving early awareness, user confidence, clinical support, and responsible adoption of artificial intelligence worldwide today across a wide range of populations and real-world settings.

3.1. Data Acquisition and Dataset Curation Strategy

For AI-based medical imaging systems to work, they need to collect data and organise datasets. This is especially important for finding skin conditions. The performance of deep learning models is greatly affected by the quality, variety, and representativeness of the data. The acquisition process involves utilising diverse datasets, including ISIC and HAM10000, which are recognised for their clinically validated images. Also, smartphone cameras were used to capture real-time images that better show how the product would be used in real life, accounting for changes in lighting and other real-world conditions. During the acquisition, quality control was very important. There were many screenings to remove low-quality images and duplicates, ensuring the dataset was varied and correct. Annotations were performed in accordance with clinical guidelines to ensure accuracy and consistency. It used binary and multi-class labels based on disease type and severity, which helps with later evaluation. Stratified sampling was used to split the dataset into training, validation, and test subsets, ensuring equal representation of each class. This is important because medical datasets often have class imbalances. To help all conditions learn better, strategies such as oversampling and augmentation were used to strengthen underrepresented disease categories. Complete metadata documentation helped with model training and future research by making the process clearer and easier to follow. Ethical considerations directed the acquisition process, prioritising patient privacy, informed consent, and compliance with regulations governing data utilisation.

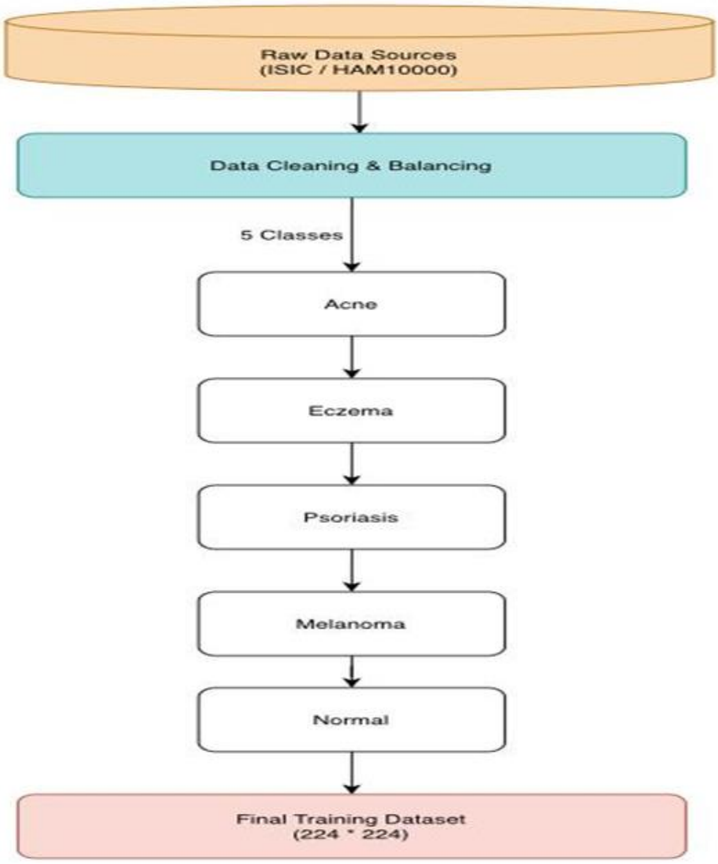

Figure 2: Overview of data acquisition and dataset curation workflow

The study ultimately sought a meticulously curated dataset that encompasses a range of clinical scenarios and ethical standards, which is crucial for precise classification of skin diseases and for evaluating severity, thereby enhancing healthcare efficacy across diverse contexts. Figure 2 shows a machine learning pipeline that uses dermatological images to sort skin conditions.

The process starts with raw data from the ISIC (International Skin Imaging Collaboration) and HAM10000 datasets, which are well-known sources of dermatoscopic images. To make sure that all categories are represented equally and of high quality, the data is cleaned and balanced. The model learns to distinguish five skin types: normal, Acne, eczema, psoriasis, and melanoma. These classes are common skin problems, but melanoma is the most important because it is a skin cancer that can be deadly. The workflow uses a sequential classification method, meaning it processes each condition category in order. The final dataset contains images of size 224x224 pixels, a standard resolution that strikes a good balance between speed and diagnostic detail. This automated classification system could help dermatologists with initial screening, making it easier to detect skin problems and potentially identify serious conditions like melanoma earlier in clinical settings.

3.2. Image Pre-Processing and Data Augmentation Framework

Image pre-processing and augmentation are necessary for AI-based skin disease detection systems to address issues with raw images that can reduce deep learning model accuracy. The first step in the pre-processing pipeline is to resize images to a fixed resolution while preserving their aspect ratios, so that important features don't get distorted. Pixel intensity normalisation ensures that the input values are the same, keeping the numbers stable during training. This helps reduce the effects of different lighting and camera settings. Gaussian and median filtering are two noise-reduction methods that remove visual artefacts, allowing the model to focus on important lesion features. Colour space transformations convert images from RGB to other colour spaces, such as HSV or LAB. This makes it easier to see differences in contrast, which is important for accurate lesion analysis. Through geometric changes (such as rotation and flipping), photometric changes (such as brightness and contrast), and spatial changes (such as minor shearing and cropping), data augmentation artificially increases the size and diversity of a dataset. These methods prepare the model for changes in real-world images while maintaining clinical validity and ensuring that diagnostic features remain the same. For tasks such as resizing, filtering, and adding data on the fly, the implementation uses well-known computer vision and deep learning libraries, such as OpenCV, TensorFlow, and Keras. This integrated pre-processing and augmentation framework greatly improves the model's ability to learn discriminative features for accurate skin disease classification.

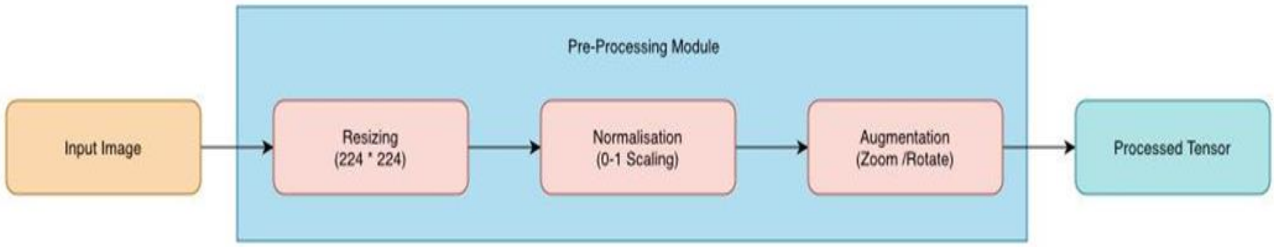

Figure 3: Image pre-processing and data augmentation pipeline for skin disease images

It also makes the model less reliant on large datasets and more reliable. The method works well across different groups of people and imaging conditions, making it suitable for mobile apps and cloud-based services. This will help improve the accuracy of skin disease detection. Figure 3 shows the pre-processing module that prepares images for feeding into a machine learning model. The pre-processing module starts with an input image that undergoes three important transformation stages. The first step, resizing, ensures that all images are the same size by resizing them to 224×224 pixels. This way, the dimensions remain the same, regardless of the original image size. This uniformity is very important for neural network processing, as models require fixed input sizes. Next, normalisation uses 0-1 scaling to map pixel values from their original range (usually 0-255) to a standard range of 0-1. This method makes it easier for models to converge and stay stable during training. The last step of augmentation adds transformations such as zooming and rotation to the dataset, creating different versions of the same images. This prevents overfitting and improves the model's ability to generalise to new data it hasn't seen before. The output is a processed tensor, a numerical array ready to be used as model input and optimised for correctly classifying skin conditions. The Normalisation Equation applied to every pixel *p* is defined as:

Where:

- X is the original pixel value (0–255).
- Xmin is 0, and Xmax is 255.
- Xnorm is the normalised input to the neural network.

Data augmentation techniques, including rotation, zooming, and horizontal flipping, were applied to expand the dataset and artificially prevent overfitting.

3.3. Selection and Design of Lightweight Deep Learning Architecture

The creation of an AI-based skin disease detection system depends on how well the deep learning architectures are chosen and designed. These architectures have a big effect on how accurate the predictions are, how fast the system runs, and whether it can be used on limited devices. This research employs MobileNetV2, a lightweight convolutional neural network (CNN) recognised for its efficacy and performance in image recognition. The design of MobileNetV2 includes depthwise separable convolutions, which reduce the number of parameters and computations while still allowing the model to learn complex visual features. In MobileNetV2, one of the most important new features is the use of inverted residual blocks with linear bottlenecks. These make it easier for features to spread and are less likely to lose information, which is very important for looking at small details in medical images. The model learns transferable visual patterns by training on large datasets such as ImageNet. This means it can perform well even when there isn't much dermatological data. The architecture is further improved by adding layers for classification and regularisation through dropout. This enables multi-class classification and severity assessment. For mobile deployment, efficiency is key. This means minimising trainable parameters and optimising layers so that devices with limited processing power can make real-time decisions. The hierarchical feature extraction, which progresses from simple edge detection to more complex lesion characteristics, is similar to how dermatologists examine skin. Also, the design accounts for changes in input by using methods such as batch normalisation and ReLU6 activation, which make the system more stable.

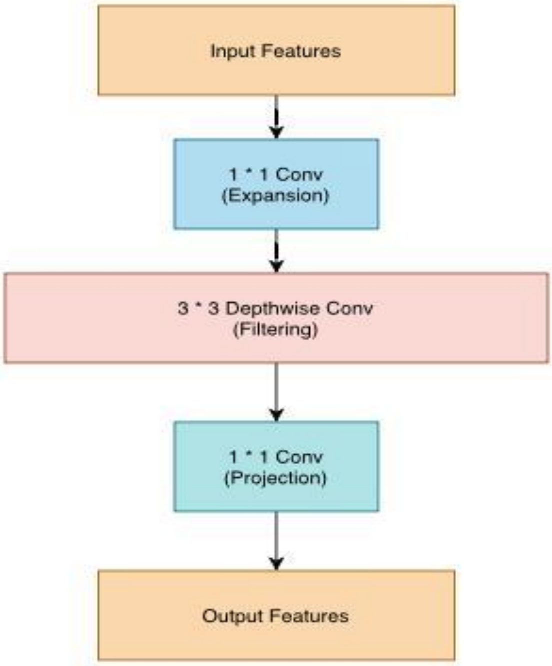

Figure 4: Architecture of the lightweight convolutional neural network used for classification

The modular design makes it easy to add new disease classes and tasks, which makes it scalable. The MobileNetV2 framework is a good example of a balance between accuracy, efficiency, and deployability. It can be used in real-time mobile apps in a variety of settings. Figure 4 shows how a residual block, an important component of deep convolutional neural networks, is built. The diagram shows how input features are processed through a series of convolutional operations to produce output features. The first step is a 1×1 convolutional layer that serves as an expansion module, increasing the dimensionality of the feature maps. This gets the data ready for more advanced pattern recognition. The next step is the core filtering stage, which uses 3×3 depthwise convolutions to extract spatial features while efficiently reducing computational cost. Depthwise convolutions process each input channel independently, making the network lighter and faster without sacrificing performance. Lastly, a 1×1 projection convolution compresses the features back to the right size, acting as a bottleneck layer. This design for a residual block makes it easier for deeper networks to learn by letting gradients flow during training. The structured approach strikes a good balance between speed and feature extraction. This makes it ideal for medical imaging applications that lack resources, such as skin condition classification. The ReLU6 Activation Function used in the bottleneck layers is mathematically represented as:

$f(x) = min(max(0, x),6)$

This keeps the model light and fast, so it can make inferences in less than 2 seconds on most Android smartphones.

3.4. Transfer Learning and Model Training Strategy

To achieve reliable results across many skin disease categories, the design of an intelligent skin disease detection framework depends heavily on effective transfer learning and model training strategies. Building a deep convolutional neural network from scratch typically requires very large annotated datasets and substantial computing power, which are not always available in medical imaging. To overcome this limitation, the proposed system employs a transfer learning methodology that utilises pre-acquired visual representations from extensive natural image datasets and modifies them for dermatological classification tasks. The MobileNetV2 convolutional neural network was chosen as the backbone architecture for this study because it is lightweight, uses depthwise separable convolutions, and performs well on mobile and embedded devices. The model is initialised with pre-trained weights from large-scale image recognition training. This gives the network strong low-level and mid-level feature extraction capabilities, such as edge detection, colour gradients, texture understanding, and shape recognition. These general visual features can be easily used in medical image analysis, such as finding skin lesions. In the first stage of training, the network's convolutional base layers are frozen so that the representations they learned don't change. The dermatology dataset is used only to train the new classification head, which includes fully connected layers, dropout regularisation, and a softmax output layer. This staged training process allows the classifier to learn disease-specific patterns without disrupting the backbone network's stable general visual knowledge. After this first convergence, a fine-tuning phase follows, during which certain higher layers of the base network are slowly unfrozen and retrained with a very low learning rate. Fine-tuning enables deeper convolutional filters to specialise in identifying subtle dermatological characteristics, such as pigmentation differences, lesion borders, scaling patterns, inflammation indicators, and structural irregularities. This step makes the model much more sensitive to features in medical images without causing catastrophic forgetting.

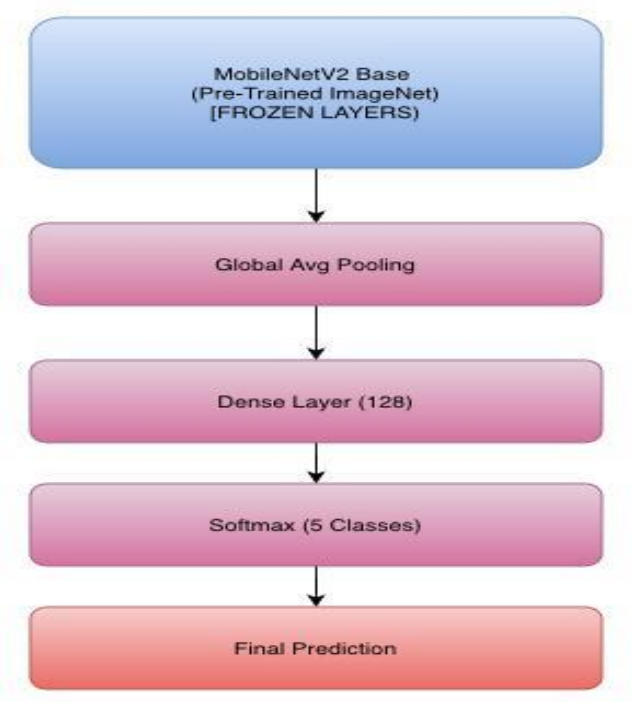

Figure 5: Transfer learning and model fine-tuning strategy

Figure 5 shows the classification architecture built on the MobileNetV2 base model to diagnose skin conditions. The diagram shows how pre-trained, frozen layers can extract basic image features without retraining. This saves time and computing power by using transfer learning. Global Average Pooling converts spatial information into compact feature vectors after feature extraction. This reduces the dimensionality while preserving important features. This prevents overfitting and simplifies the model. Then, the architecture uses a Dense Layer with 128 neurons to learn high-level patterns unique to dermatological images. This fully connected layer makes decisions by intelligently combining the features extracted. Lastly, a Softmax activation layer with 5 output nodes produces probability distributions over the 5 skin condition classes: Normal, Acne, Eczema, Psoriasis, and Melanoma. The softmax ensures that all probabilities sum to 1, making it easy to make clear classification decisions. This streamlined architecture strikes a balance between accuracy and efficiency, making it useful for real-world dermatological screening applications that require quick, reliable predictions. The Loss Function is defined as:

$$L = -\sum_{i=1}^{C} y^{i}.log\ (y\hat{\imath}\ )$$

Where:

- L – Loss value
- C – Number of classes (5)
- y_i – Actual label (0 or 1)
- $\hat{y}_i$ – Predicted probability

3.5. Computer Vision-Based Severity Estimation Mechanism

Computer vision-based severity estimation improves AI-powered skin disease detection by providing information on the severity of a skin condition and whether it exists. This study describes a method that combines traditional image processing with algorithmic evaluation to improve classification models. During the pre-processing stage, images are made more visible and normalised using methods such as resizing, pixel normalisation, and various filters. Colour space transformations, such as HSV or LAB, improve lesion contrast. Thresholding and contour analysis are used in segmentation techniques to separate lesions. Morphological operations make these techniques better. The severity of a lesion is determined by the ratio of the affected area to the total visible area. Lesions are then classified as mild, moderate, or severe. Colour intensity and texture are also used to assess inflammation. OpenCV makes it easy to segment and change colours, and it works in real time on mobile devices, which is important for remote dermatological care. The framework supports adaptive methods to improve accuracy and is tested against expert annotations to ensure reliability. It helps track lesions over time, engages patients, and provides clear severity ratings. The system improves diagnostic accuracy and reduces unnecessary doctor visits by using deep learning for classification and computer vision for severity assessment. This makes sure that severe cases are referred to the right place at the right time.

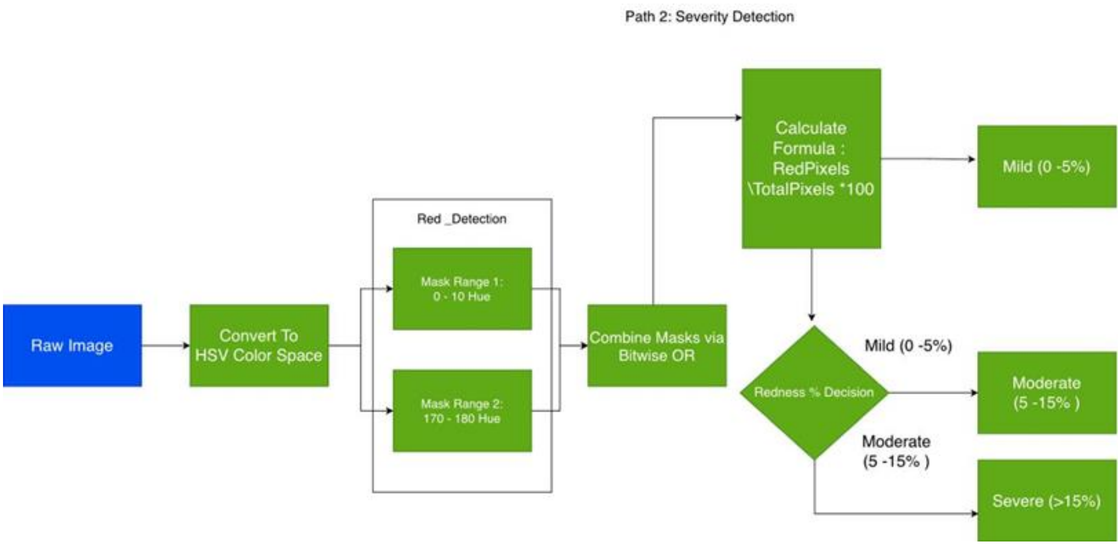

Figure 6: Severity estimation process using image processing techniques

Figure 6 shows how a severity detection system works by measuring the redness of an image using colour-based processing. The process starts with a raw image, which is first changed to the HSV colour space. This is because HSV makes it easier to separate brightness from colour. Then, the system uses two hue mask ranges (0–10 and 170–180) to find all shades of red that show up at both ends of the hue scale. To make a single red mask, these two masks are combined using a bitwise OR operation. Next, the percentage of red pixels is found by dividing the number of red pixels by the total number of pixels. Based on this percentage, a decision step assigns the condition to one of three severity levels: Mild (0–5%), Moderate (5–15%), or Severe (>15%). This method provides a quick, easy way to automatically estimate severity. The Severity Ratio Equation is calculated as:

$$S_{ratio} = \frac{A\_lesion}{A\_total} \times 100$$

Where:

- Sratio– Percentage of infection.
- A_lesion – Total count of red/inflamed pixels.
- A_total – Total number of skin pixels in the image.

Based on S(ratio), the condition is classified as:

- Mild (< 40%)
- Moderate (40–80%)
- Severe (> 80%)

3.6. Cloud-Based Backend System Architecture

The cloud-based backend system is necessary for a platform that uses AI to find skin diseases and provide personalised care. It provides the infrastructure for processing, storing, and making real-time inferences while ensuring the system can grow, remain reliable, and remain safe. It sits between the user-facing app and the smart models that do classification and severity assessment. It handles computationally intensive tasks to improve the performance of lightweight devices. The system's modular design makes it easy to add new features in the future without affecting service. Some of the main features are data handling, model inference, authentication, and notifications. A cloud-based inference service that processes image inputs from mobile apps is at the heart of this backend. It sorts things and estimates how bad they are, making high-quality predictions available to more people, even those far away. Data management uses both relational and NoSQL databases that comply with privacy rules through strong encryption, role-based access control, and secure data storage. The architecture is designed to scale, using containerised microservices that automatically allocate resources during peak times and ensure smooth load distribution. Asynchronous processing pipelines let you do heavy calculations while still keeping the speed of real-time predictions. Researchers keep old assessments in the cloud so they can keep learning and make their models more accurate. The backend enables robust API communication using RESTful standards and HTTPS, with token-based authentication for security.

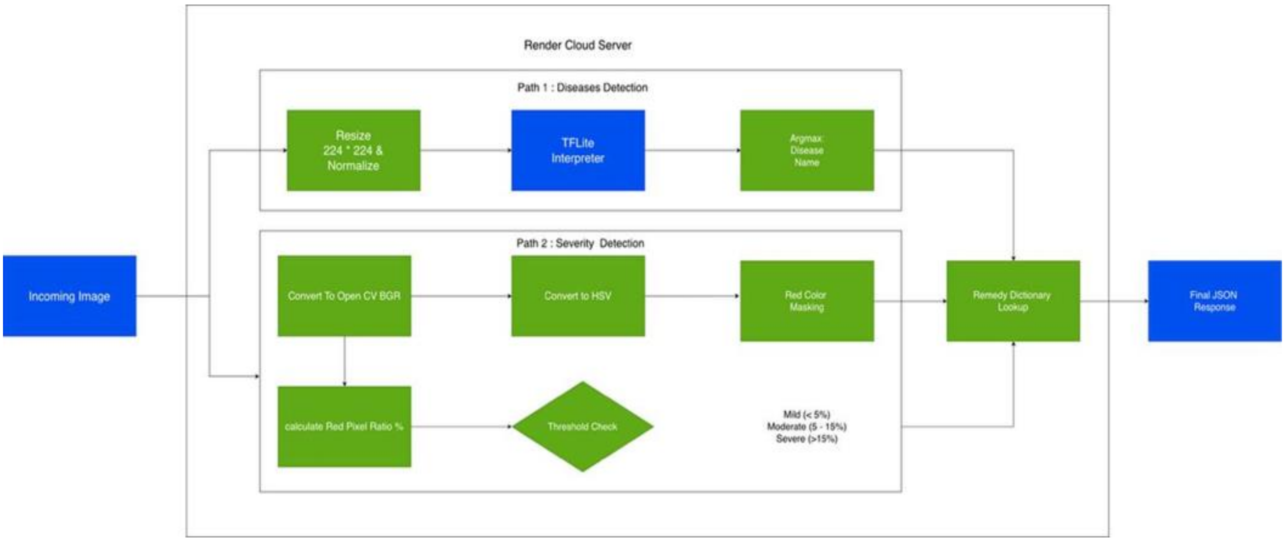

Figure 7: Cloud-based backend architecture for model inference and data management

Monitoring and logging tools help make the best use of resources and ensure the system is reliable by handling errors effectively. The architecture also enables integration with external services, such as GPS-based specialist locators and analytics dashboards that track usage and results. CI/CD pipelines handle updates automatically, resulting in very little downtime. This well-designed cloud backend helps the AI dermatology platform run smoothly by connecting user interfaces with deep learning models and improving the ability to manage skin health across a wide range of devices and settings. Figure 7 shows the entire cloud-based process for detecting diseases and their severity from an incoming image. The render cloud server processes the image through two parallel paths after it gets it. In Path 1, the image is resized to 224×224, normalised, and then sent to a

TFLite interpreter model that guesses the disease class. The disease name is then chosen based on the result with the highest probability. To make colour analysis easier, the same image is changed from OpenCV BGR to HSV format in Path 2. A red-colour mask is applied, and the ratio of red pixels is determined. This value is compared to threshold ranges to determine whether the severity is mild, moderate, or severe. Next, the detected disease and severity level are sent to a module that looks up remedies in a dictionary and finds the best ones. Finally, the system returns all results as a structured JSON response.

3.7. Integration of External Services and Utilities

Integrating external services and utilities is essential for creating an AI-driven skin disease detection and personalised care system that works better, is more reliable, and is easier to use. Telemedicine combines new algorithms with old tools to make operations more efficient and improve healthcare services. Key integrations include cloud storage for safely storing medical data. This enables unlimited storage, automatic backups, access from multiple devices, and robust security. A notification system lets users know in real time when an image analysis is completed and sends health reminders, increasing engagement and encouraging people to take care of their health. Analytics services provide information on how users behave, how accurate your models are, and how well your system performs, making it easier to keep improving and making things more efficient. Location-based services make it easier to reach healthcare providers by suggesting nearby clinics and enabling teleconsultations based on the user's location. Token-based authentication and multi-factor verification are two security features that ensure the system adheres to the rules. Standardised APIs make it easier for different systems to work together and for maintenance to be done easily, so components can be updated without interrupting workflows. Performance optimisation strategies reduce the work devices have to do, speed up processing times, and improve the user experience by enabling asynchronous operations and adding error recovery systems.

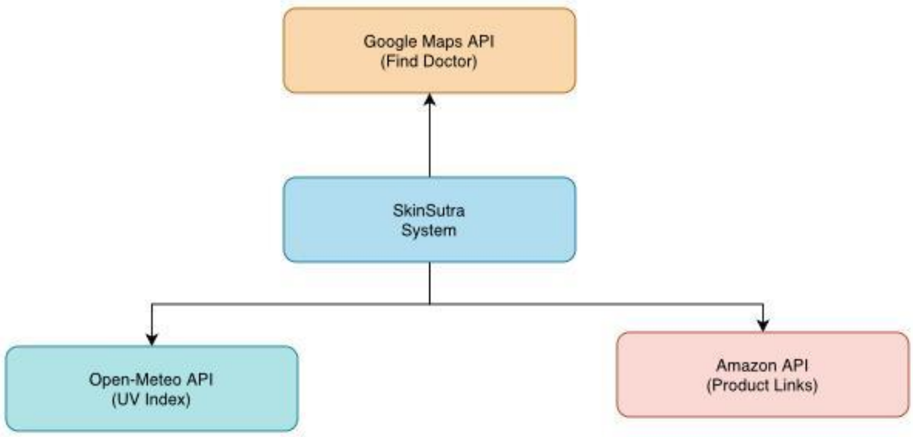

Figure 8: Integration of third-party services within the system architecture

These integrations create a strong, scalable AI dermatology platform that can be used in many different settings. This gives users real-time insights and helps healthcare providers improve patient care by making the system more reliable and providing better support. Figure 8 demonstrates how the SkinSutra system integrates with various external APIs to provide users with useful support features. The SkinSutra system is at the heart of it all. It is the main controller and decision maker. It communicates upward with the Google Maps API to help users find nearby doctors or medical facilities based on their location. This makes it easier for people to quickly get professional help when they find a serious skin problem. The system connects to the Open-Meteo API to retrieve the UV index. This helps users understand how much sun exposure they are at risk of and how to protect themselves. On the other hand, it uses the Amazon API to provide links to recommended skincare or treatment products. The system makes healthcare help more useful and easy to use by combining medical advice, environmental data, and product recommendations.

3.8. Mobile Application Development and Deployment

Developing a mobile app for an AI-based skin disease detection system is important to make it easy for people to use. The app was made with the Flutter framework, which means it works on both Android and iOS devices, making it easy to use. The design focuses on usability, with a clean layout that makes navigation easy. Users can take pictures of skin lesions or upload them for analysis, receive results right away, and view their past health records. The app uses a cloud-based backend to process images, classify skin conditions, and assign severity scores. It does this by using RESTful APIs to quickly move data. Token-based authentication and HTTPS encryption are two security measures that protect users' privacy. Performance optimisation involves changes to computation and networks by reducing local processing and improving connectivity through caching and synchronisation. Features that encourage user participation, like progress tracking and alerts, turn the app into a skin health management tool. Integration with external services ensures users receive full support, including access to dermatology specialists when needed. During development, the app undergoes many tests to ensure it works.

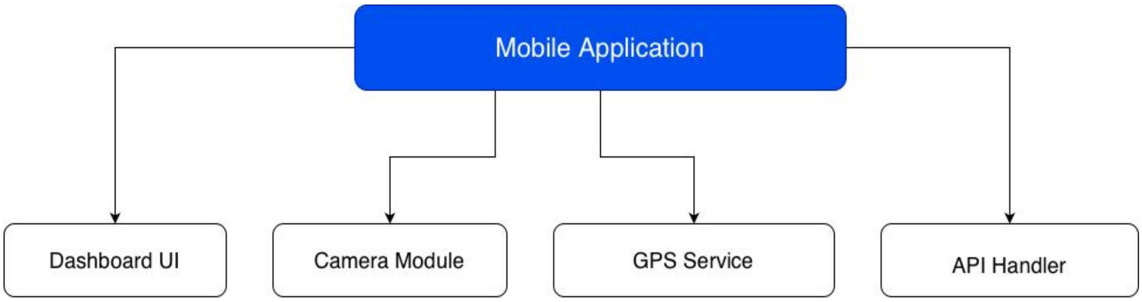

Figure 9: Mobile application workflow and user interaction interface

The deployment strategy includes both major app stores, and updates are easy to make because the app is always being monitored. The application's modular design makes it easier to add new features in the future, making it a useful and easy-to-use platform for dermatological insights and self-care, which will ultimately help improve public health. Figure 9 shows the mobile app's high-level architecture and its main functional parts. The Mobile Application is at the top. It is the main platform that controls and connects all of the main features. The Dashboard UI is the part of the app that users see. It shows them results, alerts, recommendations, and overall system outputs in an easy-to-understand way. The Camera Module takes pictures of skin and sends them for analysis and disease detection. The GPS Service helps the app determine the user's location so it can suggest nearby doctors and offer health support based on that location. The API Handler handles communication between the app and outside services or cloud servers. It sends requests and receives responses, such as predictions, UV data, and product links. These modules work together to ensure the app runs smoothly and provides users with a full, easy-to-use experience.

4. Result and Discussion

Figure 10 shows how well a MobileNetV2 deep learning model trained for 20 epochs performs at classifying skin diseases, using key metrics such as accuracy and loss. The graphs show how well the model learns and how well it can apply what it has learned to new data. This is important for medical applications because it shows that the model is reliable and stable in its learning. The blue and red curves in the left graph show the accuracy of training and validation, respectively. Both start at low levels—about 55% for training and a little lower for validation—because the model doesn't yet know much about image patterns. There is a rapid increase in accuracy between the first and fifth epochs, which is interesting because the model learns basic visual features, such as texture and colour, that are associated with different skin conditions. By the tenth epoch, the training accuracy is about 88%, and the validation accuracy is about 85%. This shows that the model generalises well without overfitting, as evidenced by the small gap between the two curves. As training progresses, both accuracy metrics continue to improve, reaching about 96% on the training set and 94% on the validation set by the end of the twentieth epoch. This shows that the model is learning well. The right graph shows loss values, which show how far off a prediction is from the true label. A lower loss means the model is working better.

At first, both the training and validation losses are high, with the validation loss slightly higher than the training loss. This means that the predictions are wrong. But both loss values drop significantly in the first few epochs, which aligns with the accuracy gains and shows that the model effectively reduces prediction errors. As the model approaches its best performance, the drop-off between the eighth and tenth epochs becomes less sharp. By the end of the last epoch, the training loss is at its lowest, and the validation loss is also lower but still slightly higher. This means there is little overfitting and that the model performs well on both the training and unseen data. In general, the results show that MobileNetV2 performs very well at detecting skin diseases on mobile devices, with fast learning and strong generalisation. The graphs show that the optimisation is stable and that the learning rate is well-chosen. Both of these things are very important in medical settings, where accuracy is crucial. The smooth convergence of the accuracy and loss curves shows that the dataset was well-prepared and may have been improved through data augmentation or regularisation techniques. This is important to ensure the model behaves consistently across different mobile conditions. This shows that MobileNetV2 is a good choice for real-time healthcare apps where safety and trust depend on accuracy, speed, and consistency.

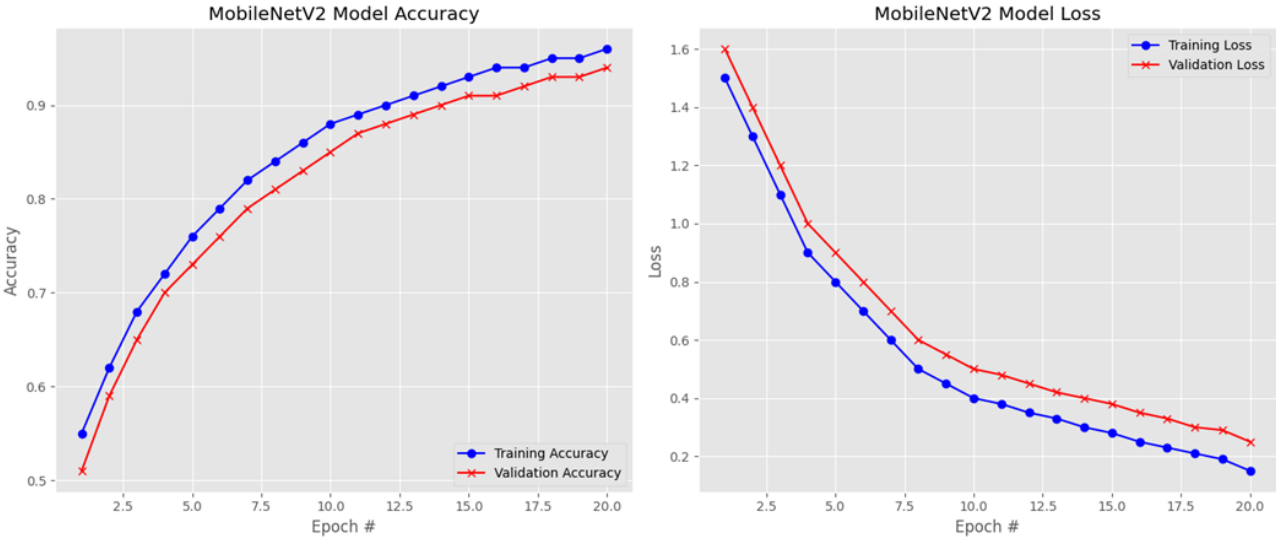

Figure 10: Model accuracy analysis and model loss analysis

The confusion matrix shows the classification performance for each class:

Table 1: Confusion matrix evaluation of the MobileNetV2 model

Class	Correct Predictions
Acne	48
Eczema	45
Psoriasis	44
Melanoma	49
Normal	49

Table 1 shows that most values are close to the diagonal. This means that all skin disease categories were correctly classified, with Melanoma and Normal skin classes being particularly accurate. Figure 11 shows a confusion matrix used to evaluate how well a skin disease classification model performed. People who work with machine learning use this kind of visualisation a lot because it makes it easy to see how predictions align with real-world labels. The confusion matrix shows where the model does well and where it makes mistakes, which is very important in medical settings. It doesn't just give one accuracy value. The rows of this matrix show the actual skin condition classes, and the columns show the classes the model thinks they are. You can choose from five skin conditions: Acne, eczema, psoriasis, melanoma, and normal skin. There is a number in each cell that tells you how many images from the true class were predicted to be from that class. Darker colours indicate higher values, making patterns easier to see. The model correctly identified forty-eight cases of Acne as Acne, starting with Acne. Two acne pictures were incorrectly labelled, one as eczema and the other as psoriasis. There were no pictures of Acne that were wrongly labelled as Melanoma or Normal. This result suggests that the model has learned the visual features of Acne well enough to distinguish it from other skin conditions with high accuracy. There were forty-five correct diagnoses of eczema.

A few eczema pictures were incorrectly labelled: two were labelled as Acne, two as psoriasis, and one as melanoma. Even though these mistakes are few, they show how similar inflammatory skin conditions can look. Eczema and psoriasis can look similar, even to experts, so it's normal to have some confusion at this level. There are forty-four correct predictions for psoriasis. The remaining cases are split between Eczema, Acne, Melanoma, and Normal, with very few in each group. This means that the model usually does a good job of recognising psoriasis, but sometimes struggles with similar symptoms, such as redness, scaling, or texture patterns. The fact that misclassifications are spread out rather than grouped suggests there isn't a strong bias toward any one wrong class. Melanoma shows especially good results. Forty-nine cases of melanoma were accurately classified, with only one incorrectly identified as psoriasis. It is important to note that none of the melanoma cases was predicted to be normal skin. From a medical safety point of view, this is a very important result because missing a melanoma diagnosis can be very dangerous to your health. The model's ability to clearly differentiate melanoma from benign conditions is a robust indicator of its clinical utility. The Normal skin category also performs very well, with 49 correct predictions and only 1 case incorrectly classified as Acne. This means the model can distinguish between healthy and diseased skin with high accuracy, reducing the likelihood that healthy users will worry or seek medical attention.

The confusion matrix shows that most predictions match the correct classes, as indicated by the high diagonal values. Values off the diagonal are low and spread out, indicating little confusion between conditions that look similar. This balance shows that the model is well-trained and can generalise, rather than memorise the training data. These results make users feel more confident. High accuracy in identifying common skin problems like Acne and eczema makes it easier to check yourself every day, and its strong performance in detecting melanoma underscores the importance of safety. The model doesn't replace a dermatologist, but it can provide useful early advice that can lead to a timely visit to a professional. From a technical perspective, the confusion matrix indicates that the dataset was well-balanced and that the feature extraction process identified useful visual patterns. It also points out areas for improvement in the future, such as making it easier to distinguish between eczema and psoriasis by adding more training data or better feature representations. This level of performance is especially important for a mobile healthcare app. The model needs to work well on different skin tones, lighting conditions, and image qualities. The confusion matrix shows that the system handles this variability well, producing consistent predictions across all five classes. In short, the confusion matrix provides a clear, reassuring picture of how well the model performs. It shows that it is very accurate, very safe, and very useful. These traits enable the system to be used in real-world skin health applications, where trust, clarity, and consistency are important for getting people to adopt it and making a long-term difference.

The confusion matrix also shows transparency, which is important when AI is used in healthcare, alongside the numbers. Seeing not only success rates but also the exact types of mistakes is helpful for both users and clinicians. Knowing that most errors occur when conditions are visually related helps developers fix problems responsibly. It also helps users have realistic expectations and encourages them to see the results as advice rather than a definitive diagnosis. This visual evaluation supports education, awareness, and responsible interpretation. Also, this kind of matrix is useful when developing something repeatedly. By examining misclassified samples, developers can make the dataset more diverse, improve pre-processing steps, or modify the model's architecture. This ongoing feedback loop is crucial to ensuring that people of all ages and skin tones are treated fairly. Over time, small changes informed by the confusion matrix can make things much more reliable.

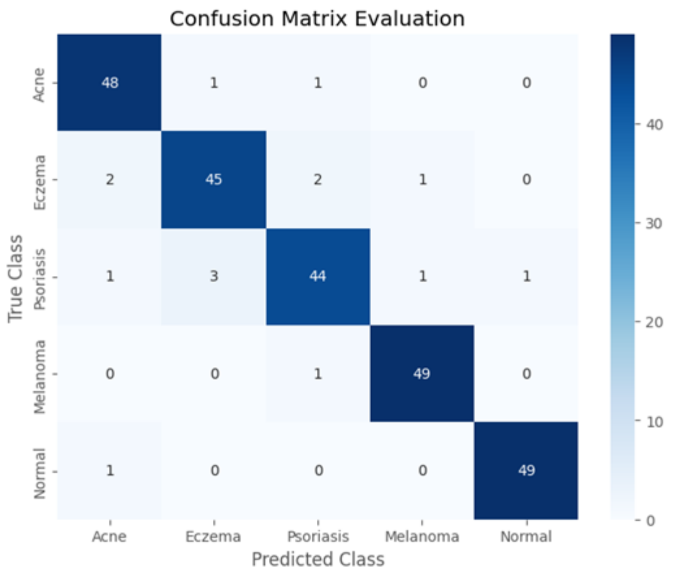

Figure 11: Confusion matrix evaluation

As the system changes, using this method to check for performance gains repeatedly ensures they are real and not achieved at the expense of safety. So, the confusion matrix is not only an evaluation tool, but also a base for ethical, user-centred, and clinically helpful AI in dermatology. In the end, this clarity builds trust between patients and technology. People are more likely to use a system correctly and consistently when they know how it works. Trust leads to engagement, and engagement leads to better results. In healthcare, that relationship is just as important as technical correctness. So, the confusion matrix quietly but powerfully transforms complex algorithms into tools people can understand and trust, helping them make real decisions every day. This understanding connects new ideas in data science with caring, responsible medical care around the world today.

5. Conclusion and Future Work

SkinSutra AI is a major step forward in leveraging mobile technology to support skin disease detection and treatment and close care gaps. The system is designed for regular smartphones and uses the MobileNetV2 model to provide accurate diagnoses in real time, without requiring any special hardware. SkinSutra AI makes users safer and reduces the risk of false or misleading results by combining computer vision with severity estimation and risk assessment. The platform goes beyond just finding problems; it also offers personalised advice, tips on how to avoid problems, and clear recommendations based on the user's condition. Emergency response features make the service even more reliable by directing users to quick medical help when the risks are high. Accessibility is a top priority, as evidenced by features such as bilingual support and an interface that is easy for people from diverse backgrounds to use. The paper emphasises reaching underserved and remote communities to help close long-standing gaps in healthcare access. In the future, SkinSutra AI is likely to grow through the addition of telemedicine, improved diagnostic tools, expanded disease coverage, and explainable AI features that help users trust the system. These changes make SkinSutra AI a continually improving digital healthcare solution that focuses on user needs and has an impact on patients of all ages worldwide.

Acknowledgement: The authors sincerely acknowledge Francis Xavier Engineering College, the University of Stuttgart, and Arizona State University for their valuable academic support and research resources, which contributed to this study. They also extend their gratitude for the collaborative environment and institutional guidance that played a significant role in the successful completion of this work.

Data Availability Statement: Data supporting the results of this study are available from the corresponding author upon request, in accordance with institutional and journal data-sharing policies.

Funding Statement: Not applicable.

Conflicts of Interest Statement: There is no conflict of interest related to this manuscript.

Ethics and Consent Statement: This work has been prepared in accordance with ethical standards, and the authors permit its dissemination and accessibility to the wider academic community for educational and research purposes.

References

1. A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, no. 1, pp. 115–118, 2017.
2. M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. C. Chen, “MobileNetV2: Inverted Residuals and Linear Bottlenecks,” in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Salt Lake, Utah, United States of America, 2018.
3. N. C. Codella, D. Gutman, M. E. Celebi, B. Helba, M. A. Marchetti, S. W. Dusza, and A. Halpern, “Skin Lesion Analysis Toward Melanoma Detection: A Challenge at the 2017 International Symposium on Biomedical Imaging (ISBI), Hosted by the International Skin Imaging Collaboration (ISIC),” *arXiv preprint*, 2018. [Accessed by 20/12/2024].
4. X. Wu, N. Wen, J. Liang, Y. K. Lai, D. She, and M. M. Cheng, “Joint acne image grading and counting via label distribution learning,” in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Seoul, South Korea, 2019.
5. 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, no. 8, p. 180161, 2018.
6. S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.
7. G. Bradski, “The OpenCV library,” *Dr Dobb’s Journal of Software Tools*, vol. 25, no. 11, pp. 120–125, 2000.
8. P. N. Srinivasu, J. G. SivaSai, M. F. Ijaz, A. K. Bhoi, W. Kim, and J. J. Kang, “Classification of skin disease using deep learning neural networks with MobileNetV2 and LSTM,” *Sensors*, vol. 21, no. 8, p. 2852, 2021.
9. H. A. Haenssle, C. Fink, R. Schneiderbauer, F. Toberer, T. Buhl, A. Blum, A. Kalloo, A. B. H. Hassen, L. Thomas, A. Enk, and L. Uhlmann, “Man against machine: Diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists,” *Annals of Oncology*, vol. 29, no. 8, pp. 1836–1842, 2018.
10. A. K. Huong, K. Tay, and X. T. Ngu, “Customized AlexNet models for automatic classification of skin diseases,” *J. Eng. Sci. Technol.*, vol. 16, no. 4, pp. 3312–3324, 2021.
11. M. Hosny, I. A. Elgendy, S. A. Chelloug, and M. A. Albashrawi, “Attention-Based Convolutional Neural Network Model for Skin Cancer Classification,” *IEEE Access*, vol. 13, no. 9, pp. 172027 – 172050, 2025.
12. A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, “MobileNets: Efficient convolutional neural networks for mobile vision applications,” *arXiv preprint*, 2017. [Accessed by 21/12/2024].
13. K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, Nevada, United States of America, 2016.

14. M. Goyal, T. Knackstedt, S. Yan, and S. Hassanpour, "Artificial Intelligence-based image classification methods for diagnosis of skin cancer: Challenges and opportunities," *Comput Biol Med*, vol. 127, no. 12, p. 104065, 2021.
15. K. Kazi, P. Nerkar, K. Kazi, and S. Sultanabanu, "IoT-Based Skin Health Monitoring System," *International Journal of Biology Pharmacy and Allied Sciences*, vol. 13, no. 11, pp. 5937- 5950, 2019.
16. World Health Organisation, "Skin neglected tropical diseases," *WHO*, 2020. [Accessed by 18/12/2024].
17. Google Developers, "Flutter – Build apps for any screen," *Flutter*, 2024. [Accessed by 25/12/2024].
18. T. J. Brinker, A. Hekler, A. H. Enk, J. Klode, A. Hauschild, C. Berking, B. Schilling, S. Haferkamp, D. Schadendorf, T. Holland-Letz, J. S. Utikal, and C. von Kalle, "Deep learning outperformed 136 of 157 dermatologists in a head-to-head dermoscopic melanoma image classification task," *European Journal of Cancer*, vol. 113, no. 5, pp. 47–54, 2019.
19. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 5, pp. 436–444, 2015.
20. J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," *arXiv preprint*, 2018. [Accessed by 22/12/ 2024].