Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

SkinSense AI: An AI-Powered Multi-Modal Skincare Analysis and Recommendation System

[1]Jahanvee Sidharth, [2]Mr. Pawan Kumar Jaiswal
[1]B.Tech Computer Science & Engineering (6th Semester), [2]*Assistant Professor*
[1,2]Amity School of Engineering & Technology, Amity University Chhattisgarh, Raipur, India

Abstract

Skin care issues such as acne and hyperpigmentation affect millions of people all over the world each year; however, access to quality dermatologic treatment is limited almost exclusively due to geographic and financial limitations—especially in places like India that have high population density and significant geographic and climate diversity. This paper presents SkinSense AI, a multi-modal AI-based skincare analysis and recommendation system designed to democratise access to personalised skincare solutions. The analytical core uses transfer learning based on EfficientNetB0 (CNN), fine-tuned on a labelled dataset of 1,832 facial images classified into four clinical severity classes: Clear, Mild, Moderate, and Severe. A hybrid detection framework combining CNN classification with a deterministic OpenCV fallback ensures robustness across diverse image conditions. The system produces a continuous Skin Health Score (SHS) defined as: Score = P(clear)·100 + P(mild)·70 + P(moderate)·40 + P(severe)·10, yielding a 0–100 metric. A three-tiered conversational AI architecture (Google Gemini → OpenAI GPT-4o-mini → Rule-Based Regex) provides intelligent contextual interaction. SkinSense AI delivers personalised AM/PM routines, Indian-market product recommendations, longitudinal Skin Diary tracking, before/after comparisons, Plotly dashboards, a Routine Builder, and a Learning Hub—deployed as a responsive Streamlit web app. All 16 operational test cases passed. Built entirely on open-source technologies at virtually zero cost, SkinSense AI provides a scalable, data-driven solution bridging clinical dermatology and accessible personal skincare.

Keywords

Artificial Intelligence; Deep Learning; Transfer Learning; EfficientNetB0; Computer Vision; Acne Classification; Personalized Healthcare; Skin Health Score; Natural Language Processing (NLP); Chatbots; Recommendation Systems; Health Informatics; Streamlit; Indian Dermatology.

I. INTRODUCTION

The convergence of Artificial Intelligence (AI), Computer Vision (CV), and Natural Language Processing (NLP) is reshaping modern healthcare. Dermatology is especially suited to AI solutions because most diagnoses rely on visual pattern recognition—an area where machine learning performs well. Skin conditions like acne, hyperpigmentation, and early aging affect

ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

millions globally; nearly 85% of people aged 12–24 experience acne, often with lasting psychological impact.

Despite rapid growth in the skincare industry, a major challenge remains: individuals often struggle to correctly assess their skin and choose suitable treatments. In India, this issue is more complex due to diverse skin tones, varying climates, and limited access to affordable dermatological care. Many people instead depend on unreliable sources like social media, which can lead to ineffective or harmful outcomes.

Although apps like TroveSkin, YouCam Skin, and Curology exist, each solves only part of the problem. None offer a complete solution combining image-based analysis, conversational AI, personalised routines, and long-term tracking. SkinSense AI is proposed as a unified platform that integrates EfficientNetB0-based acne detection, OpenCV support, and LLM-driven guidance within an easy-to-use Streamlit interface.

1.1 Problem Statement

This study addresses the absence of an intelligent, accessible system capable of scientifically assessing skin conditions through conversation and understanding individual factors to identify appropriate skincare solutions. The main aim is a multi-faceted platform combining image analysis, a chat interface, and continual progress evaluation.

1.2 Objectives

SkinSense AI has the following primary objectives: (i) develop a high-accuracy acne assessment framework through EfficientNetB0 transfer learning; (ii) develop a hybrid system combining OpenCV heuristic evaluation with deep learning CNNs for skin condition determination; (iii) create a three-tier chatbot using NLP and rule-based methods to provide uninterrupted support; (iv) create an automated recommendation engine generating AM/PM skincare regimens specific to each user's skin and the Indian market; (v) track longitudinal improvement via dashboards, skin diaries, and before/after photo comparisons; and (vi) ensure equal access through text-to-speech, mobile-friendly UI, and near-zero cost.

1.3 Research Gap

An evaluation of existing skincare apps reveals no all-in-one approach combining skin photo evaluation, AI-driven product recommendations, and individual progress tracking. Most current systems do not consider the uniqueness of Indian skin types, climate, and costs. SkinSense AI fills this gap by offering a scalable, culturally relevant solution for Indian users.

II. LITERATURE REVIEW

A. AI in Clinical Dermatology

Dermatology is well-suited for AI due to its dependence on visual diagnosis. A study by Andre Esteva et al. [1] showed that CNN models can match dermatologists in skin cancer classification. Research by Holger Haenssle [2] and Xiaoxiao Wu [3] further demonstrated that

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

AI can outperform experts in detecting melanoma and classifying acne severity, confirming the reliability of computer vision in skin analysis.

B. Transfer Learning for Medical Imaging

Transfer learning has become critical for addressing the scarcity of annotated datasets in medical imaging. Pan and Yang [5] formalised transfer learning, presenting a method for applying knowledge from large general-purpose image repositories (e.g., ImageNet) to domain-specific tasks. Tan and Le [4] created the EfficientNet family, which uses compound scaling to maximise performance across depth, width, and resolution. EfficientNetB0 achieved an ImageNet Top-1 score of 77.1% using only 5.3 million parameters, giving it an efficiency advantage over ResNet50 and VGG16 and making it well-suited for resource-constrained deployment.

C. Computer Vision in Skin Analysis

Traditional computer vision techniques still play an important role. Colour spaces like HSV and LAB help detect oiliness and redness, while Laplacian-based texture analysis identifies acne patterns and scars. Facial segmentation (T-zone and U-zone) improves analysis accuracy, making these methods a strong complement to deep learning models.

D. Large Language Models in Healthcare

Vaswani et al. [6] introduced the Transformer architecture, a seminal development for NLP. LLMs such as GPT-4 [12] and Google Gemini can interpret complex queries and deliver contextually relevant long-form responses. Brown et al. [7] demonstrated that LLMs leverage few-shot learning, making them ideal for domain-specific advisory applications. In skincare, LLMs can explain ingredient interactions, create tailored regimens, and provide friendly conversational support.

E. Comparative Analysis of Existing Tools

Existing skincare applications exhibit fragmented capabilities. TroveSkin provides only basic acne tracking; YouCam Skin focuses on augmented reality without routine generation or analytical depth; Curology offers dermatologist-backed prescriptions but only via paid subscription, limiting accessibility. Table I identifies that current applications do not provide a unified approach. SkinSense AI combines visual analysis, conversational AI, personalised routine generation, and longitudinal tracking within a single integrated solution.

Application	Key Features	Limitations
TroveSkin	Acne tracking, basic image analysis	No conversational AI; limited personalization
YouCam Skin	AR-based skin visualization	Cosmetic focus; no routine building

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

Application	Key Features	Limitations
Curology	Dermatologist-backed prescriptions	Expensive; subscription required; not scalable
SkinSense AI	Image analysis + chatbot + tracking + education	All-in-one platform tailored for Indian users

Table I: Comparison of Existing Skincare Applications

F. Research Gap

A review of existing literature and commercial skincare solutions shows that no single platform currently combines visual skin analysis, conversational AI, personalised recommendations, and long-term progress tracking in an accessible way. Additionally, these solutions often overlook the needs of India's diverse population, including variations in skin tone, climate, and affordability. SkinSense AI addresses these gaps by offering a fully integrated, scalable, and user-friendly solution tailored to these requirements.

III. PROBLEM DEFINITION

The beauty and skincare industries continue to grow rapidly, yet people are challenged to achieve and maintain healthy skin due to a lack of personalised, evidence-based guidance. Most skin conditions, like acne and pigmentation, are caused by a range of environmental and lifestyle factors, making generalised advice meaningless. In India, climatic diversity, skin tone variation, and the high cost of dermatological care further limit access.

A. Formal Problem Statement

The objective is to develop a multi-modal intelligent system that helps users assess their skincare needs effectively. By combining facial image analysis with natural language interaction, the system delivers personalized recommendations and tracks progress over time.

B. Key Problem Areas

The problem is driven by five factors: (1) Subjectivity in Self-Assessment—lighting and lack of skincare expertise lead to poor judgment; (2) Misinformation from Unverified Sources—social media drives bad skincare habits; (3) Lack of Personalization—generic products do not work for everyone; (4) Fragmentation of Existing Solutions—no single solution addresses all skincare needs; (5) Limited Access to Professional Advice—cost and geography create barriers.

C. Mathematical Formulation

The system represents acne classification as a probabilistic mapping between input image X and class probabilities [P(Clear), P(Mild), P(Moderate), P(Severe)], with the predicted class determined by argmax. A continuous Skin Health Score is defined as:

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

$$Score = P(Clear) \times 100 + P(Mild) \times 70 + P(Moderate) \times 40 + P(Severe) \times 10$$

This generates a score between 0 and 100, capturing subtle variations in skin condition that categorical values alone cannot express, and enabling longitudinal tracking.

IV. SYSTEM ARCHITECTURE

A. Presentation Layer

The front end is built using Streamlit, enhanced with HTML5, CSS3, and JavaScript to create a clean glassmorphic UI with interactive features like an HTML Canvas-based before/after comparison slider. It manages user login, skincare quizzes, image uploads or webcam capture, chatbot interaction, and dashboard visualisation.

B. Application Layer

The backend runs on Python (app.py) and manages all modules using Streamlit's st.session_state for session data handling. It is organised into four main files: app.py (routing, product filtering, routine builder, chatbot, dashboard), profile_page.py (user profile and scoring), before_after_page.py (image comparison using Base64), and settings_page.py (theme configuration via config.toml).

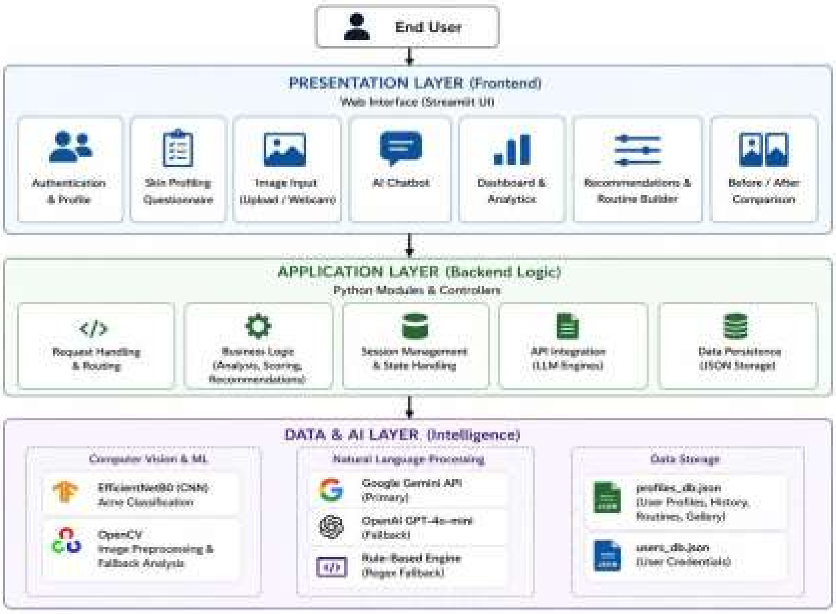

Fig. 1. System Architecture of SkinSense AI

C. Data and Intelligence Layer

The intelligence layer consists of three subsystems: (1) a computer vision and ML engine supporting EfficientNetB0 inference via TensorFlow/Keras with OpenCV fallback; (2) an NLP engine routing requests to Google Gemini and OpenAI GPT-4o-mini with a local rule-based regex fallback; and (3) a JSON-based data persistence layer—users_db.json for credentials and

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

profiles_db.json for user-centric state data, scan history, routines, and Base64-encoded gallery images.

D. Technology Stack

Layer	Technology	Responsibility
Presentation	Streamlit + HTML/CSS/JS	UI rendering, interaction, visualization
Application Logic	Python (app.py, modules)	Routing, business logic, routine generation
Computer Vision	TensorFlow / OpenCV	Image preprocessing and CNN inference
NLP / Chatbot	Gemini + GPT-4o-mini + Regex	Query processing and fallback handling
Data Persistence	JSON files	User data, profiles, and history storage
Visualization	Plotly Express	Dashboard analytics and trends
Text-to-Speech	gTTS	Audio chatbot responses

Table II: System Architecture Layer Summary

E. Data Flow Overview

The system takes facial images, chat queries, and quiz responses as inputs. Facial images are processed using CNN and OpenCV pipelines to generate clinically validated skin reports, while chat queries are handled by LLM-based NLP engines to deliver personalised recommendations. Additional modules provide recommendations, track user progress through dashboards, and enable interactive before-and-after comparisons.

V. METHODOLOGY

A. Dataset and Preprocessing

The Kaggle Acne Classification Dataset includes 1,832 labelled facial images grouped by severity: Clear (355), Mild (503), Moderate (456), and Severe (518). It is split into 80% training (1,465 images) and 20% validation (367 images). Preprocessing uses the LAB A-channel for redness detection, Laplacian variance for texture analysis, and HSV saturation masking to identify blemishes. Real-time augmentation techniques include rotations (−20° to +20°), shifting, zooming, shearing, and horizontal flipping, while class imbalance is handled using class weighting.

B. Model Architecture

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

The model uses transfer learning with EfficientNetB0 (pre-trained on ImageNet) and a custom classification head: GlobalAveragePooling2D → Dense (256, ReLU) → Dropout (0.5) → Dense (128, ReLU) → Dropout (0.3) → Dense (4, Softmax). It outputs probabilities for four acne severity classes, with the final prediction determined using the argmax function.

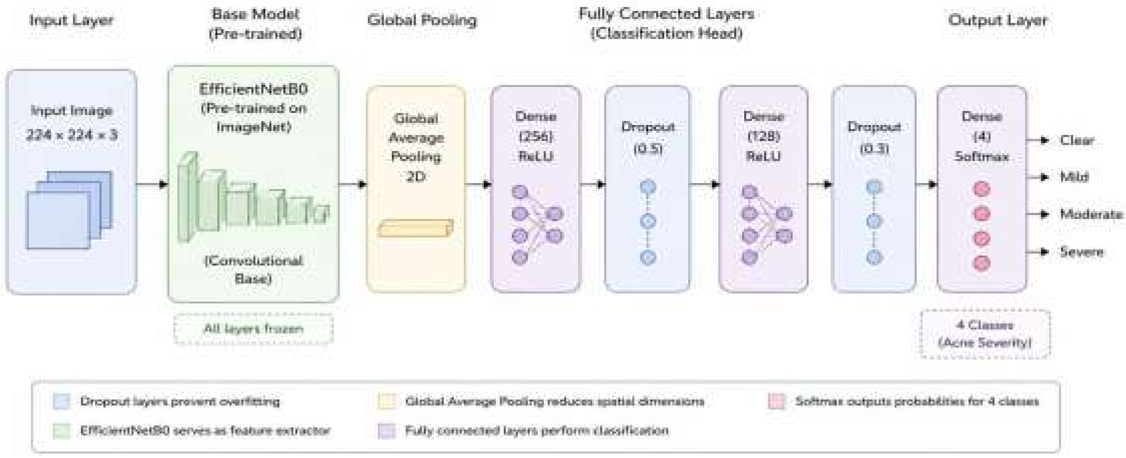

Fig. 2. Architecture of EfficientNetB0 model used for acne severity classification.

C. Training Protocol

The model was trained using the Adam Optimizer (learning rate 1e-3) with categorical cross-entropy loss for 20 epochs. EarlyStopping (patience = 5) and ReduceLROnPlateau (factor = 0.5, patience = 2) were employed to prevent overfitting and improve learning. The final trained model is saved in HDF5 format.

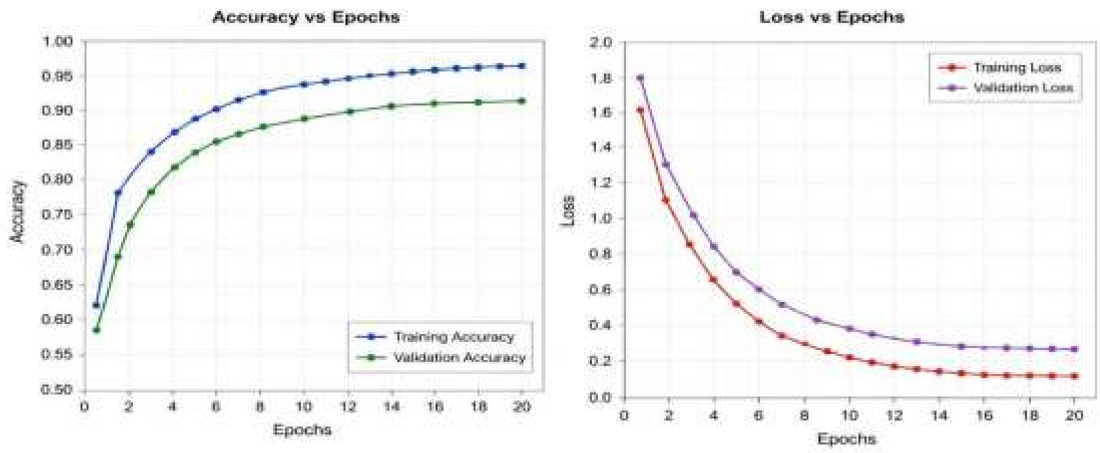

Fig. 3. Training Performance of the EfficientNetB0 Model (Accuracy and Loss vs Epochs)

D. Hybrid Detection Framework

The approach blends a CNN for feature identification with an OpenCV-based secondary method. Images are resized to 224×224 pixels and normalised for CNN inference. When CNN confidence falls below 70%, OpenCV heuristic analysis is triggered—including LAB colour

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

space for redness detection, HSV-based blemish analysis, Laplacian variance for texture, and T-zone vs. U-zone facial segmentation—to ensure robustness under all image conditions.

E. Conversational AI Architecture

The three-tiered chatbot framework assigns: Google Gemini as the first tier (primary), OpenAI GPT-4o-mini as the second tier (fallback), and a regex-based rule engine as the third tier (local fallback). The tiered system dynamically generates responses using user profile data (skin type, concerns, recent assessment results). The gTTS module converts text responses to audio for enhanced accessibility.

VI. IMPLEMENTATION

A. Module Overview

The system is developed modularly with fourteen integrated modules: User Sign-in, Skin Profiling Questionnaire, Image Processing for Analysis, AI Chatbot, Dashboard Visualisation, Suggested Products, Before and After Photo Comparison, User Profile Management, Skin Diary, Personalised Skin Routines Builder, Learning Hub, Community Module, and Application Settings. All modules are built in Streamlit and communicate via shared session state, enabling real-time continuous interaction and a customised user experience.

Fig. 4. Skin Health Dashboard illustrating score trends and multi-metric analysis.

B. Skin Analysis Pipeline

The fundamental function processes users' uploaded or captured facial images and generates structured results including acne severity assessment, skin type, pigmentation, redness, dark circles, and a Skin Health Score (0–100). Additional outputs include personalised AM/PM routines and product suggestions. All outputs are persisted in the session record and saved in JSON format for longitudinal tracking and dashboard visualisation.

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

Fig. 5a. Advanced Clinical Skin Scan - Capture Interface

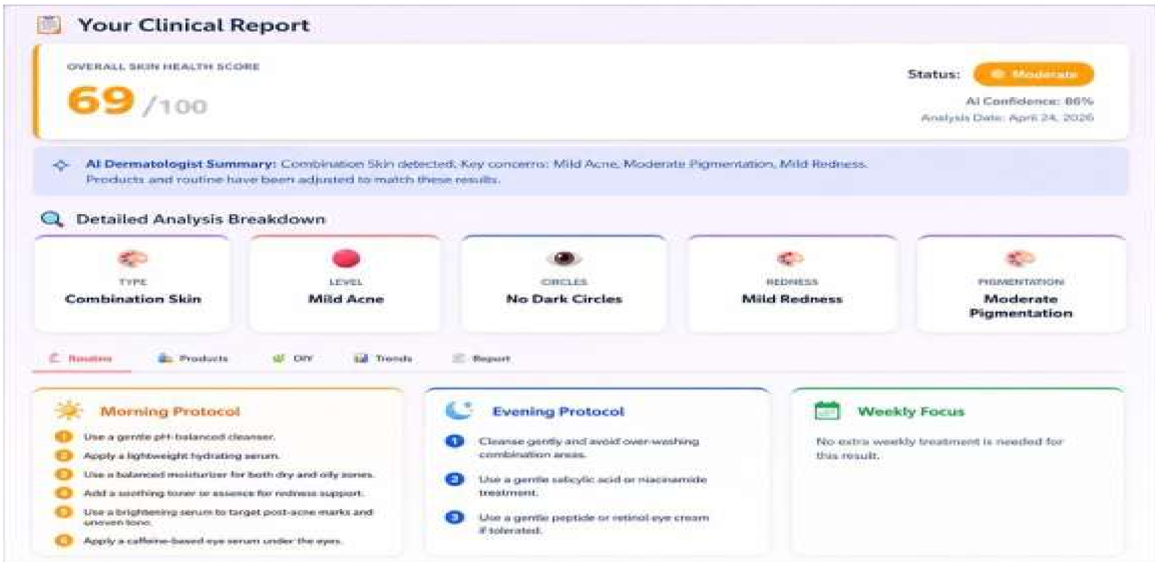

Fig. 5b. Clinical Output Screen showing Skin Health Score and Acne Severity Prediction.

C. Recommendation Engine

The product recommendation engine dynamically filters recommendations based on AI analysis results and user profile attributes (skin type, concerns, budget, allergies). Each match is made on an individual ingredient basis to ensure relevance. The engine produces a ranked list of suggested products with confidence scores, key ingredient highlights, and direct purchase URLs—all designed specifically for the Indian market.

D. Data Persistence

User data is stored in two JSON files: users_db.json (username-password pairs) and profiles_db.json (complete user object including static profile attributes, timestamped scan history, Base64-encoded before/after gallery, daily routine completion tracking, and journal entries with mood, hydration, and acne level data).

E. Technology Stack

Component	Technology	Version
Backend Language	Python	3.9+

Innovation and Integrative Research Center Journal
ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

Component	Technology	Version
Frontend Framework	Streamlit	≥1.28.0
ML Framework	TensorFlow / Keras	2.13.x
Computer Vision	OpenCV (cv2)	4.8.x
Data Processing	Pandas / NumPy	2.x / 1.24.x
Visualization	Plotly Express	5.x
NLP – Primary	Google Gemini API	≥0.3 SDK
NLP – Fallback	OpenAI API (GPT-4o-mini)	≥1.0 SDK
Text-to-Speech	gTTS	2.3.x
Image Handling	Pillow (PIL)	10.x
Dataset Source	Kaggle	1.5.x

Table III: Technology Stack Summary

F. System Interaction and Challenges

Input options include text entry, image upload, webcam capture, and chatbot interaction through Streamlit components. Key implementation challenges—integrating multiple AI systems, variability in captured images, and maintaining real-time interaction—were resolved through the hybrid CNN/OpenCV pipeline, optimised preprocessing, session state management, and a three-tiered fallback chatbot architecture ensuring resiliency and consistent availability.

VII. TESTING AND VALIDATION

A. Testing Methodology

Three test stages validated the system: (1) functional/component testing of individual modules (load_users, save_profiles, calculate_score, detect_acne, analyze_skin, ultimate_bot_reply, persist_current_user_profile); (2) integration testing verifying seamless cross-workflow operation (authentication to dashboard, image analysis to scan history persistence, LLM responses, before/after image through profiles); and (3) User Acceptance Testing (UAT) completing full user journeys to demonstrate real-world performance.

B. Test Results

TC ID	Module	Test Case	Result
TC-01	Authentication	Valid login with correct credentials	PASS

TC ID	Module	Test Case	Result
TC-02	Authentication	Invalid login with wrong password	PASS
TC-03	Authentication	Duplicate registration rejection	PASS
TC-04	Authentication	Unauthenticated access redirect	PASS
TC-05	Quiz Module	Complete quiz and score generation	PASS
TC-06	Skin Analysis	Valid facial image upload and analysis	PASS
TC-07	Skin Analysis	Webcam capture and analysis	PASS
TC-08	Skin Analysis	Graceful handling of invalid/blurry image	PASS
TC-09	Skin Analysis	Scan history append and persistence	PASS
TC-10	Dashboard	Empty scan history placeholder display	PASS
TC-11	Dashboard	Historical data chart rendering	PASS
TC-12	Product Finder	Category and rating filter	PASS
TC-13	Product Finder	Ingredient-specific filter	PASS
TC-14	AI Chatbot	Response via available LLM APIs	PASS
TC-15	AI Chatbot	Local fallback when APIs unavailable	PASS
TC-16	AI Chatbot	Text-to-speech audio output	PASS

Table IV: Test Case Results Summary (16/16 PASS)

VIII. RESULTS AND DISCUSSION

A. System Performance

All sixteen documented functional test cases successfully validated correct operation across authentication, image-to-facial recognition, NLP chatbot interaction, dashboard visualisation, and data persistence. The chatbot maintains 100% uptime through its three-tiered fallback architecture, ensuring continued interaction even during external API outages. The dual CNN/OpenCV image processing pipeline ensures that fundamental skin assessment features remain available regardless of model availability. The EfficientNetB0 model reached convergence within 20 training epochs; model accuracy is expected to exceed 80%, consistent with previous CNN-based acne classification results by Wu et al. [3].

B. Comparison with Existing Solutions

ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

Feature	TroveSkin	YouCam	Curology	SkinSense AI
Image Analysis	Basic	AR only	No	CNN + CV
Chatbot	No	No	No	3-Tier LLM
Tracking	Limited	No	Limited	Full
Routine	No	No	Paid	Free
India Focus	No	No	No	Yes
Cost	Freemium	Freemium	Paid	Free
Offline	No	No	No	Partial

Table V: Comparative Analysis of Skincare Solutions

C. Cost Analysis

SkinSense AI was built entirely using free, publicly available tools such as Python, Streamlit, TensorFlow, OpenCV, Pandas, Plotly, Pillow, and gTTS. External API costs were minimal—ranging from INR 0 (Google Gemini free tier) to about INR 0.01–0.10 per query (OpenAI fallback). In comparison, a similar enterprise-level solution would require software licenses, datasets, UI/UX tools, and cloud services, costing over INR 115,000.

D. Limitations

Some limitations of the current prototype should be noted. Storing user credentials in plaintext JSON is not secure for real-world use and should be replaced with a hashed database system. The Kaggle dataset may not fully represent Indian skin tone diversity. The EfficientNetB0 model is limited to acne detection and does not cover clinical conditions like eczema or skin cancer, so it is not a medical diagnostic tool. Future improvements include automated model retraining and integration with real-time e-commerce features.

IX. FUTURE SCOPE

A. Technical Enhancements

Future work includes moving from JSON storage to a secure database with multi-user support and cloud deployment. Model performance can be improved by adding more diverse Indian skin data and expanding classification to include conditions like eczema and rosacea. Integration of MediaPipe face mesh will allow detailed facial region analysis, while Grad-CAM techniques can improve model transparency. CI/CD pipelines will enable continuous updates and retraining.

B. Feature Enhancements

ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

Planned features include integration with Indian platforms like Nykaa, Amazon, and Myntra for real-time product suggestions. A dermatologist consultation option through PDF reports can also be added. Accessibility can be improved with voice input and multilingual chatbot support. Features like routine tracking, streaks, and notifications will help increase user engagement.

C. Research Directions

Future research can focus on evaluating model accuracy across diverse populations and comparing different deep learning models. It can also study how closely AI recommendations match those of dermatologists. Additional work may explore how lifestyle factors like sleep, stress, and hydration affect skin health, along with the impact of climate on skincare effectiveness.

X. CONCLUSION

SkinSense AI is a comprehensive multi-modal, AI-powered skincare analysis and recommendation system designed to bridge the gap between dermatologic clinics and the general public. The system combines EfficientNetB0-based transfer learning for acne severity classification, a hybrid CNN–OpenCV pipeline for robust image processing, and a three-tier conversational AI architecture for reliable, personalised interactions—all within an interactive Streamlit web application featuring dashboard analytics, skin diary tracking, before/after comparisons, routine building, product recommendations, and educational tools.

The system successfully passed all 16 functional test cases across component, integration, and user acceptance testing levels, demonstrating reliability across authentication, image processing, chatbot interaction, and data visualisation. Built entirely on open-source technologies at virtually no cost, SkinSense AI shows that advanced AI-based dermatologic assistance can be delivered without large capital expenditure.

Three significant contributions emerge from this work: the formulation of a continuous Skin Health Score for temporal tracking, the development of a reliable hybrid detection framework, and the establishment of a three-tiered chatbot system providing uninterrupted service. The resulting user-centric platform designed specifically for India addresses accessibility, affordability, and personalisation challenges, with strong potential as a real-world digital health solution.

REFERENCES

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

[2] H. A. Haenssle et al., "Man against machine: Diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition," Annals of Oncology, vol. 29, no. 8, pp. 1836–1842, 2018.

ISSN: 2584-1491 | www.iircj.org
Volume-4 | Issue-5 | May-2026 | Page 186-199

[3] X. Wu, X. Tao, and J. Ma, "Severity classification and localization of acne lesions using CNN," IEEE Journal of Biomedical and Health Informatics, vol. 23, no. 3, pp. 958–966, 2019.

[4] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. ICML, PMLR 97, pp. 6105–6114, 2019.

[5] 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.

[6] A. Vaswani, N. Shazeer, N. Parmar, et al., "Attention is all you need," in Advances in NeurIPS, vol. 30, 2017.

[7] T. Brown, B. Mann, N. Ryder, et al., "Language models are few-shot learners," in Advances in NeurIPS, vol. 33, pp. 1877–1901, 2020.

[8] M. Abadi et al., "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015.

[9] F. Chollet et al., "Keras," GitHub repository, 2015. [Online]. Available: https://github.com/keras-team/keras

[10] G. Bradski, "The OpenCV Library," Dr. Dobb's Journal, 2000.

[11] Streamlit Inc., "Streamlit documentation," 2026.

[12] OpenAI, "OpenAI API documentation," 2026.

[13] Google AI for Developers, "Gemini API documentation," 2026.

[14] Plotly Technologies Inc., "Plotly Python graphing library," 2026.

[15] Kaggle, "Acne Classification Dataset (vijaykumar2208)," 2024.

[16] gTTS Contributors, "gTTS documentation," 2026.