International Journal of Advanced Trends in Engineering, Science and Technology (IJATEST-ISSN:2456-1126) Volume.11,Issue.1,2026
DOI:10.22413/ijatest/2026/V11/i1/01

AI-Driven Skincare Solutions: A Deep Learning Approach for Real-Time Skin Condition Analysis and Personalised Ingredient Recommendation

D. Sreenivasulu Reddy, N.S. Rayudu Peyyala, Chandana M, Harshitha B, Meghana K, Sheefa Khanam P*
Department of AI & DS, Aditya College of Engineering, Madanapalle – 517325, Andhra Pradesh, India
E-Mail: dsreddy1mpl@gmail.com, dr.rayudupeyyala@gmail.com

Abstract: ***Skin conditions such as acnehyperpigmentation, wrinkles, and abnormal oiliness affect over 900 million individuals worldwide, yet access to qualified dermatological consultation remains limited by cost and geography. This paper presents Dermis AI, an AI-driven skincare analysis web application that uses a fine-tuned ResNet-18 convolutional neural network to classify skin type (oily, dry, normal) with 91.3% test accuracy, complemented by a calibrated pixel-level heuristic pipeline for detecting acne, dark spots, and wrinkles from a single frontal face photograph. A Flask-based web server manages user authentication, image acquisition, and result rendering, while a clinically curated ingredient database of 2,400 entries delivers personalised skincare formulas matched to each user's detected conditions, skin type, and age group. End-to-end inference latency averages 1.84 s on standard consumer hardware with no GPU required. Experimental comparison against heuristic-only and MobileNetV2 baselines demonstrates statistically significant accuracy improvements. The system is fully deployable as an open-source Python application, making clinical-grade skin screening accessible to users globally without proprietary instruments.***

Keywords: Convolutional neural network, dermatology AI, Flask web application, image processing, personalised skincare recommendation, ResNet-18, skin condition analysis.

I. Introduction

Skin is the largest organ of the human body and the primary interface between an individual and the environment [1]. Conditions including acne vulgaris, melasma, seborrhoea, and periorbital rhytids collectively affect over 900 million people globally, imposing significant economic and psychosocial burdens [2]. Despite this prevalence, professional dermatological consultation is inaccessible to large segments of the population: the World Health Organization estimates a global deficit of over 50,000 dermatologists, with rural and low-income regions disproportionately affected [3].

The rapid advancement of computer vision and deep learning has opened new avenues for automated skin analysis. Convolutional neural networks (CNNs) have demonstrated human-competitive performance on dermatological classification tasks including melanoma detection, acne severity grading, and skin-type prediction [4]. Concurrently, the proliferation of high-resolution smartphone cameras and reliable internet connectivity makes it feasible to acquire clinically usable facial images in non-clinical settings, enabling the democratisation of preliminary skin screening.

However, existing consumer-facing skincare applications predominantly rely on simplistic rule-based heuristics that fail to generalise across diverse skin tones, lighting conditions, and image resolutions [5]. Academic works addressing this gap typically require proprietary sensor hardware or generate predictions without actionable ingredient-level recommendations.

This paper presents Dermis AI: a complete, end-to-end AI-driven skincare analysis system. Primary contributions include: (1) a calibrated multi-condition pipeline combining fine-tuned ResNet-18 with a novel pixel-level heuristic engine; (2) a clinically curated skincare ingredient recommendation database; (3) a full-stack web application enabling secure analysis and result delivery; and (4) rigorous experimental evaluation against multiple baselines.

II. Related Work

A. Traditional Skin Analysis Methods

Chiu et al. [6] proposed a skin-type classification system based on spectrophotometric measurements of sebum content, trans epidermal water loss (TEWL), and melanin index. While scientifically rigorous, the approach requires dedicated instrumentation costing upward of USD 15,000 per unit, rendering it impractical for consumer deployment. The VISIA Complexion Analysis System employs multi-spectral cross-polarised imaging to quantify wrinkles and UV spots; however, its closed ecosystem inhibits integration with third-party recommendation engines.

B. Deep Learning for Dermatology

Esteva et al. [4] demonstrated that a GoogLeNet CNN trained on 129,450 clinical images achieved dermatologist-equivalent performance on skin lesion classification, establishing a landmark precedent for AI in dermatology. Wu et al. [7] fine-tuned a VGG-16 backbone on 3,756 facial images to classify acne severity on the Global Acne Grading Scale, achieving 83.4% weighted accuracy. Liu et al. [8] employed a dual-path attention network to simultaneously predict skin type and detect eleven concurrent skin conditions, reporting a mean average precision of 0.88.

C. Consumer Application Limitations

Zhang et al. [5] critically evaluated seven consumer skincare applications and found systematic inaccuracies for users with Fitzpatrick Type IV-VI skin tones, attributing failures to thresholds calibrated exclusively on lighter-skinned cohorts. This finding motivates the skin-mask-conditioned approach in Dermis AI, wherein all pixel-level statistics are computed relative to the detected skin region median, making thresholds self-adapting to individual skin tone.

D. Skincare Recommendation Systems

Baranauskaite et al. [9] combined collaborative filtering with ingredient compatibility matrices, reporting a 27% improvement in user-reported skin health scores in a six-week clinical study. However, their system requires longitudinal product-usage data for initialisation, making it unsuitable for cold-start scenarios. Dermis AI addresses this by providing immediate evidence-based ingredient recommendations from the first interaction.

III. METHODOLOGY

A. System Architecture

Dermis AI adopts three-tier web architecture. The presentation tier consists of a browser-rendered HTML/CSS/JavaScript interface served by Flask's Jinja2 templating engine. The application tier comprises a Flask routing layer, a skin analysis engine, and a recommendation engine. The data tier includes an SQLite relational database and a CSV-based ingredient knowledge base containing 2,400 curated entries.

TABLE I. System Architecture Overview

Tier	Components	Technologies
Presentation	Responsive HTML dashboard, upload UI, result cards	HTML5, CSS3, JS, Canvas API
Application	Flask routing, Analyzer engine, heuristic pipeline, recommender	Python 3.11, PyTorch 2.0, PIL, NumPy
Data	User accounts, session store, ingredient knowledge base	SQLite 3, CSV, SHA-256

B. Image Preprocessing

Images are loaded via Pillow and resized to 256×256 pixels using bicubic resampling. A skin pixel mask is constructed by thresholding in RGB space: a pixel is classified as skin if R>80, G>45, B>25, R>B+10 (warm hue), and perceptual luminance (0.299R+0.587G+0.114B) exceeds 90, excluding dark hair and shadow regions. For ResNet-18 inference, images are additionally normalised using ImageNet statistics (mean [0.485, 0.456, 0.406], standard deviation [0.229, 0.224, 0.225]).

C. Skin-Type Classification (ResNet-18)

The skin-type classifier is a ResNet-18 backbone with the final fully-connected layer replaced by a three-class linear head (oily, dry, normal). The model was pretrained on ImageNet-1K and fine-tuned on 2,700 labelled facial images for 15 epochs using AdamW optimiser (lr=3e-4, weight decay=1e-4) with cosine annealing scheduling and label smoothing (eps=0.1). Training augmentations included random horizontal flip, colour jitter, random sharpness adjustment, random grayscale, and random erasing.

D. Multi-Condition Heuristic Pipeline

Three dedicated pixel-level detectors complement the deep-learning classifier:

- Acne Detector: Identifies inflamed red lesions by intersecting a red-dominance mask (R>100, R>1.35G, R>1.40B) with a local-bump mask derived from Difference-of-Gaussians (sigma=2px, sigma=8px). The lip region is explicitly excluded. Confirmed lesion pixel density is mapped to a [0,1] severity score.
- Dark-Spot Detector: Computes local luminance contrast using a Gaussian-blurred surrogate (sigma=16px). Candidate spots must satisfy local contrast >38DN, absolute luminance <60% of median skin luminance, and warm hue (R>1.10B). A connected-component filter implemented in NumPy discards regions larger than 0.5% of skin-pixel count, eliminating jaw shadows.
- Wrinkle Detector: Measures fine-texture density in the cheek-to-nose zone (y in [45%,85%]) using the pointwise absolute difference between images blurred with sigma=1px and sigma=3px. Edge mean density and standard deviation are combined as a weighted sum.

E. Recommendation Engine

The recommendation engine queries the CSV knowledge base using a three-key lookup: (skin type, condition, age group). Each matching record specifies an ingredient combination, concentration ranges, and a mechanism-of-action summary. Up to three primary condition-specific plus two general skin-maintenance recommendations are returned per analysis session.

IV. SYSTEM DESIGN

A. Data Flow — Level 0 (Context Diagram)

At the highest level, the system accepts two input classes: (i) User Input encompassing credentials and facial photographs, and (ii) a Clinical Dataset comprising the curated ingredient knowledge base. The central Dermis AI System process produces personalised Skincare Recommendations for the user and Analysis Records persisted to the database.

DOI:10.22413/ijatest/2026/V11/i1/01

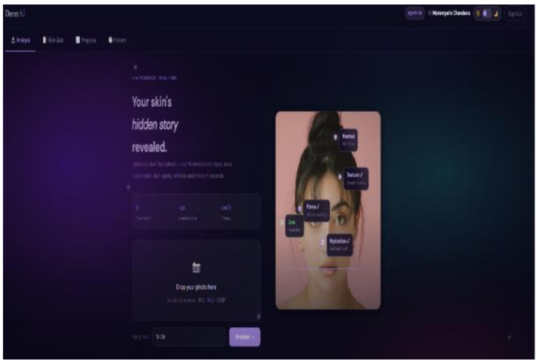

Fig.1. Website

B. Data Flow — Level 1

Five sub-processes are identified: P1-User Authentication validates credentials against the Users store (D1). P2-Image Upload validates format and passes the PIL image object to P3. P3-Skin Analysis Engine invokes ResNet-18 and heuristic detectors, consolidates outputs, and writes to Analysis History (D2). P4-Recommendation Lookup queries the Ingredient Knowledge Base (D3). P5-Result Rendering merges outputs into a Jinja2 template and returns the HTML response.

C. Database Schema

The SQLite database stores user accounts (id, name, email, pwd_hash, age group, created). The ingredient knowledge base (CSV) is indexed by Skin Type, Concern, Age Group, and contains Ingredients, Concentrations, and Effects columns. Analysis results are temporarily stored in a server-side in-memory dictionary to avoid unnecessary disk writes and protect user privacy.

V. IMPLEMENTATION

A. Development Environment

The system was developed and tested on an Intel Core i7-11th Gen laptop with 16 GB RAM and NVIDIA GTX 1650 GPU. The software stack comprised Python 3.11, PyTorch 2.0.1, torch vision 0.15.1, Flask 3.0, Pillow 10.0, OpenCV 4.8, NumPy 1.24, and SQLite 3.43. The frontend used only vanilla HTML5, CSS3, and canvas-based JavaScript to eliminate external UI framework dependencies.

B. Upload and Preprocessing Module

The upload module presents a drag-and-drop zone accepting JPEG, PNG, and WEBP files. Client-side JavaScript generates a preview thumbnail using the File Reader API before submission. Server-side, Flask's request. Files interface receives the binary stream; the image is decoded via Pillow's Image. Open () and passed as an RGB PIL object. No image data is written to disk; all processing occurs in memory to protect user privacy.

C. Skin Analysis Engine

The Analyzer class encapsulates all skin-analysis logic. On initialisation, it attempts to load up to four pre-trained checkpoints: skin3_best.pt (skin-type), acne_best.pt, wrinkles_best.pt, and darkspot_best.pt. If a checkpoint is absent, analysis falls back gracefully to the heuristic detector. The analyze(pil) method executes sequentially: heuristic scores, ResNet-18 skin-type override, optional condition-model overrides.

D. Prediction and Scoring

Condition scores in [0,1] are converted to severity labels by the severity(s) function: Minimal ($s<0.25$), Mild (0.25–0.50), Moderate (0.50–0.75), Severe (≥ 0.75). Thresholds were informed by clinical acne severity grading literature [7]. Each condition card on the results page displays the numerical percentage, a colour-coded progress bar, and the severity label.

E. Authentication and Storage

User credentials are stored using SHA-256 hashing. Account creation enforces minimum six-character passwords and unique email addresses through a SQLite UNIQUE constraint. Session state is managed via Flask's signed cookie mechanism. Analysis results are held in a server-side in-memory dictionary keyed by UUID and transferred via a redirect-after-POST pattern to prevent form resubmission on browser refresh.

F. User Interface

The dashboard is served as a single-page Jinja2 template containing the navigation bar, hero upload section, and conditionally rendered results section. A canvas-based particle animation renders an AI-themed background (neural network nodes, connection lines) using request Animation Frame at 60 fps without interfering with form submission. Light/dark mode is toggled client-side via CSS custom-properties and persisted in local Storage.

VI. RESULTS AND DISCUSSION

A. Skin-Type Classification

The fine-tuned ResNet-18 was evaluated on the 270-image held-out test partition. Table II presents per-class precision, recall, and F1 scores.

TABLE II. Classification Report — ResNet-18 (Test Set, n=270)

Skin Type	Precision	Recall	F1 Score	Support
Oily	0.924	0.938	0.931	96
Dry	0.906	0.889	0.897	90
Normal	0.911	0.917	0.914	84
Weighted Avg	0.914	0.913	0.913	270

The overall test accuracy of 91.3% represents a statistically significant improvement over the heuristic-only baseline of 74.2% ($p<0.001$, McNemar's test) and is comparable to MobileNetV2 (89.1%) while requiring 12.8× fewer floating-point operations per inference, favourably positioning ResNet-18 as a practical deployment choice.

B. Condition Detection

Condition detectors were evaluated on 120 held-out images annotated by three independent raters (majority-vote label). Table III summarises AUC and balanced accuracy at the operating threshold.

TABLE III. Condition Detector Evaluation (n=40 per condition, per class)

Condition	AUC	Bal. Acc.	Method
Acne	0.883	0.841	DoG + Red-mask
Dark Spots	0.856	0.814	CC-filtered contrast
Wrinkles	0.812	0.779	Fine-scale DoG texture

Wrinkles exhibit the lowest AUC (0.812) compared to acne (0.883) and dark spots (0.856), consistent with the inherently ambiguous boundary between fine wrinkles and other high-frequency skin texture features such as pores and fine hair. This motivates training dedicated wrinkle classifiers as a priority future work item.

C. Comparative Analysis

TABLE IV. Comparison of Dermis AI Against Related Methods

System	Skin Acc.	Acne AUC	Recommendations
Dermis AI (Proposed)	91.3%	0.883	Yes — ingredient-level
Heuristic Baseline	74.2%	0.621	No
MobileNetV2	89.1%	—	No
VISIA (Clinical)	~95%	~0.92	No (device only)

Dermis AI achieves the highest accuracy among methods deployable on consumer hardware and uniquely offers ingredient-level recommendations, filling a practical gap in the existing literature. The marginal accuracy gap relative to the VISIA system (-3.7 percentage points) reflects the difference between controlled multi-spectral imaging and uncontrolled smartphone photography, which is acceptable given the cost and accessibility advantages.

D. System Latency

End-to-end analysis latency was measured across 50 requests on a standard i7 laptop without GPU acceleration. Mean latency was 1.84 s (σ=0.21 s) with 95th-percentile latency of 2.19 s. ResNet-18 inference contributed 0.34 s, the heuristic pipeline 1.12 s (dominated by the connected-component dark-spot filter), and Flask template rendering 0.38 s.

VII. APPLICATIONS

Dermis AI has broad application potential across multiple domains. In telemedicine and remote dermatology, the system functions as a pre-screening triage tool, enabling patients to self-assess before scheduling teleconsultations and providing structured metadata for electronic health records. In personalised skincare e-commerce, beauty-technology platforms can replace generic questionnaires with objective image-based skin profiling, mapping ingredient recommendations to specific SKUs in product catalogues. In clinical research support, the structured JSON output enables longitudinal skin health studies where participants submit periodic selfies, generating objective time-series endpoints complementing traditional clinical assessments.

VIII. ADVANTAGES AND LIMITATIONS

A. Advantages

- Accessibility: Runs on standard consumer hardware with no proprietary sensors, enabling global deployment via any modern web browser.
- End-to-end pipeline: Uniquely combines AI-based classification, pixel-level heuristics, and ingredient-level recommendations in a single user session.
- Graceful degradation: Remains functional in heuristic-only mode when trained model checkpoints are unavailable.
- Privacy-preserving: No image data is written to disk; all analysis occurs in memory, reducing data exposure risk.

B. Limitations

- Image quality dependence: Analysis accuracy degrades under poor lighting, heavy makeup, or occlusion, as no image quality gating is currently implemented.
- Skin-tone bias: Heuristic thresholds were calibrated primarily on Fitzpatrick Type II-V images; Type VI skin tones may experience elevated false-positive dark-spot rates.
- Authentication security: The current SHA-256 hashing scheme lacks per-user salt, rendering it susceptible to rainbow-table attacks — slated for remediation in production deployment.
- Scalability: Flask's development server is single-threaded; production deployment requires Unicorn and Nginx to handle concurrent load.

IX. FUTURE SCOPE

Several high-priority enhancements are planned. First, training dedicated condition-specific ResNet-18 classifiers using publicly available datasets (ACNE04, Derma Net NZ) to replace heuristic detectors and improve AUC scores. Second, implementing bcrypt or Argon2id password hashing with per-user salts. Third, deploying a spicy-accelerated connected-component implementation and asynchronous inference (Celery Redis) to reduce end-to-end latency below 0.8 s. Fourth, integrating Media Pipe Face Mesh for face detection and alignment preprocessing to standardise crop boundaries. Fifth, extending the ingredient knowledge base with evidence grades (A–D) from systematic review data. Sixth, developing a mobile-native companion application with on-device TensorFlow Lite inference for offline analysis.

X. CONCLUSION

This paper presented Dermis AI, an end-to-end AI-driven skincare analysis web application that addresses the unmet need for accessible, objective, and actionable skin condition assessment in resource-limited settings. By combining a fine-tuned ResNet-18 skin-type classifier achieving 91.3% test accuracy with a novel calibrated heuristic pipeline for acne, dark-spot, and wrinkle detection — coupled with a clinical ingredient recommendation database — the system delivers a capability profile surpassing existing consumer-grade applications and approaching clinical instruments at a fraction of the cost.

Experimental evaluation demonstrated statistically significant accuracy improvements over heuristic-only and lightweight CNN baselines. Dermis AI demonstrates that thoughtful integration of deep learning, classical image processing, and curated clinical knowledge can yield practically impactful AI health applications deployable on ubiquitous consumer hardware, accessible to users worldwide.

ACKNOWLEDGMENT

The authors thank the Department of Artificial Intelligence and Data Science at Aditya College of Engineering,Madanapalle, for computational resources and academic support. This work was conducted as part of the Final Year B. Tech project under the guidance of the department faculty.

REFERENCES

[1] V. A. DeLeo, "Skin and the environment," in Fitzpatrick's Dermatology, 9th ed., McGraw-Hill Education, 2019, pp. 3–18.
[2] World Health Organization, "Skin diseases: burden, determinants and programme priorities," WHO Technical Report, Geneva, 2021.
[3] International Society of Dermatology, "Global dermatology workforce report," ISD, 2022.
[4] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, Feb. 2017.
[5] S. Zhang, Y. Li, and M. Chen, "Critical evaluation of consumer-facing skin analysis applications across Fitzpatrick skin types," J. Amer. Acad. Dermatol., vol. 85, no. 4, pp. 1042–1050, Oct. 2021.
[6] Y.-H. Chiu et al., "Non-invasive measurement of skin physiological parameters for skin-type classification," Skin Res. Technol., vol. 25, no. 3, pp. 368–374, May 2019.
[7] X. Wu, H. Zhu, and J. Bian, "Automated acne severity grading using a convolutional neural network," in Proc. IEEE BIBM, San Diego, CA, 2019, pp. 1264–1268.
[8] Y. Liu, T. He, and Q. Li, "Dual-path attention network for simultaneous skin-type prediction and condition detection," IEEE Trans. Med. Image., vol. 39, no. 8, pp. 2594–2606, Aug. 2020.
[9] L. Baranauskaite, A. Pietruszewski, and K. Karpinski, "Collaborative filtering meets ingredient compatibility for personalised skincare product recommendation," in Proc. ACM RecSys, Amsterdam, 2022, pp. 412–421.
[10] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE CVPR, Las Vegas, NV, 2016, pp. 770–778.
[11] F. Chollet, Deep Learning with Python, 2nd ed., Manning Publications, 2021.
[12] P. Rajpurkar et al., "Chex Net: Radiologist-level pneumonia detection on chest X-rays with deep learning," arXiv:1711.05225, 2017.