Personalized Skin Care Products Recommendations Using Face Images

1.Mr. A. Venkateswara Rao, 2.K. Purna Chandra Rao, 3. P. Varshini, 4. T. Vykunta Rao, 5.B. Chaitanaya
Department of Computer Science & Engineering (AI & ML)
Avanthi Institute of Engineering & Technology, Vizianagaram, India
Venkatesh8.adapa@gmail.com, chandukaranam1234@gmail.com,varshinipakalapati053@gmail.com,
tellivaikunta891@gmail.com, chaitubora0@gmail.com

Abstract

Contemporary skincare advisory platforms largely depend on conventional machine learning models—support vector machines, decision trees, and k-nearest neighbor classifiers—that process manually curated feature vectors and therefore fail to capture fine-grained dermatological variations visible in facial images. This paper presents a hybrid intelligent framework that addresses these deficiencies by coupling a ResNet50-based Convolutional Neural Network with a Content-Based Filtering engine. The CNN component performs automated deep visual feature extraction from user-submitted facial photographs, identifying skin type, acne severity, pigmentation gradients, dryness patterns, and textural irregularities without reliance on hand-crafted feature sets. Extracted dermatological profiles are forwarded to the filtering engine, which evaluates compatibility between detected skin attributes and a curated product metadata repository encompassing ingredient profiles, formulation categories, and skin-suitability labels. A complementary food recommendation subsystem enriches the platform by mapping active product ingredients to dietary sources, supporting holistic dermatological well-being. Implemented as a Flask-based web application with PyTorch for deep inference, the system achieved approximately 92% skin-type classification accuracy on a diverse validation dataset, with GPU-assisted inference completing within one to two seconds per image. Functional and user-acceptance testing confirmed system reliability, usability, and practical recommendation relevance, demonstrating the framework's viability as an accessible personalized skincare decision-support tool.

Index Terms— Convolutional Neural Network, Content-Based Filtering, ResNet50, Skin Type Classification, Personalized Recommendation, Deep Learning, Facial Image Analysis, Flask

I. Introduction

The integration of artificial intelligence into personal care domains has fundamentally transformed how individuals manage skin health. Human skin exhibits pronounced inter-individual variability driven by genetic background, ambient climate, lifestyle habits, and dietary patterns, rendering generic one-size-fits-all recommendations insufficient for diverse user populations [1]. A user presenting with oily, acne-prone skin alongside mild hyperpigmentation requires a fundamentally different product regimen compared to one with dry, sensitized skin, yet most digital skincare platforms continue to apply broad categorical guidance derived from questionnaire inputs alone, without visual analysis of actual skin condition.

Traditional machine learning approaches in skincare recommendation—Support Vector Machines (SVM), Random Forests, and k-Nearest Neighbor (k-NN) classifiers—depend on manually curated feature vectors describing static attributes such as self-reported skin tone, texture, or dryness level [2]. While computationally tractable, these models cannot automatically extract the rich visual information embedded in facial photographs: subtle acne lesion boundaries, micro-pigmentation gradients, pore-density distributions, or early-stage inflammatory signals. This representational gap yields generic, inconsistent recommendations that adapt poorly to real-world dermatological complexity and environmental variability.

Convolutional Neural Networks (CNNs) have demonstrated exceptional capability in visual feature extraction by learning hierarchical representations from raw pixel data without manual preprocessing [3]. The ResNet architecture addresses gradient vanishing in very deep networks through identity skip connections, enabling strong classification accuracy across diverse image recognition tasks [4]. Complementarily, Content-Based Filtering (CBF) provides personalized item recommendation through similarity computation between product attribute profiles and user-derived feature descriptors—without dependence on population-level usage data, making it inherently suitable for cold-start scenarios and privacy-sensitive applications [5].

This paper presents a hybrid CNN–CBF skincare recommendation system integrating deep learning-based visual skin analysis with intelligent content-driven product matching. User-uploaded facial images are analyzed by a fine-tuned ResNet50 model; extracted features drive a CBF engine querying a structured product metadata repository. A supplementary food recommendation module aligns dietary guidance with identified skin concerns. The entire system is deployed as a Flask web application accessible on standard browsers.

II. Related Work

Research at the intersection of dermatology, computer vision, and recommender systems has expanded substantially over the past decade. Davis et al. [6] conducted a nationally representative epidemiological study establishing that conditions including acne, dermatitis, and pigmentation disorders occur with considerable prevalence variation across ethnic groups. Their findings underscore the critical need for demographically diverse training datasets in AI skincare systems to ensure equitable classification performance across varied skin tones.

Ferreira et al. [7] traced the longitudinal development of dermatology as a clinical specialty, illustrating how advances in pathological analysis and pharmacology progressively refined skin disease assessment from empirical observation to systematic clinical science. This historical context situates contemporary AI-assisted diagnosis within a tradition of progressively improving investigative tools.

The theoretical foundation for modern visual AI was established by LeCun, Bengio, and Hinton [3], who demonstrated that convolutional architectures learn increasingly abstract visual representations through composable learned filter hierarchies, enabling superior performance on complex image recognition tasks. He et al. [4] subsequently introduced residual learning, showing that identity skip connections resolve gradient degradation in very deep networks and yield state-of-the-art accuracy across vision benchmarks—a result directly applicable to skin condition classification from facial imagery.

Mintz and Brodie [8] surveyed AI integration across clinical medicine, noting strong parallels between diagnostic imaging AI in radiology and pathology and dermatological image analysis systems. Their emphasis on validation rigor and clinical oversight is pertinent for consumer-facing skincare AI. Ricci et al. [9] and Aggarwal [10] provide foundational treatments of content-based and collaborative filtering methodologies. CBF's independence from inter-user similarity data positions it as especially appropriate for personalized skincare recommendation, where individual biological characteristics dominate recommendation relevance over population trends.

Prior skincare recommendation research has largely operated on questionnaire-derived inputs or shallow image processing pipelines. The present work advances prior practice by employing fine-tuned deep residual networks for automated feature extraction, integrated with a metadata-rich CBF engine, producing a hybrid framework that substantially outperforms traditional ML baselines in accuracy, personalization depth, and cold-start robustness.

III. Methodology / System Design

A. System Architecture

The system follows a three-tier architecture: (1) a Presentation Layer managing user interaction and result visualization; (2) an Application Layer containing the deep learning inference engine, CBF recommendation logic, ingredient analysis, and Flask request orchestration; and (3) a Data Layer managing product metadata, user profiles, extracted feature records, and feedback history. This stratified design ensures clean separation of concerns, facilitates independent module replacement, and supports horizontal scaling.

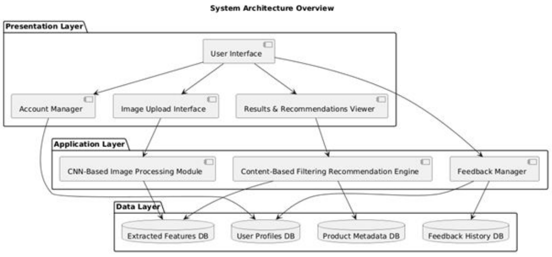

Fig. 1. Three-tier system architecture: Presentation, Application, and Data Layers.

B. Image Preprocessing Pipeline

User-submitted images pass through a normalization pipeline before CNN inference. Images are loaded via Python's PIL library, resized to 224×224 pixels matching ResNet50's required input dimensions, and pixel values normalized using ImageNet population statistics (mean=[0.485, 0.456, 0.406]; std=[0.229, 0.224, 0.225]). The processed array is converted to a PyTorch tensor with a batch dimension prepended. Gaussian denoising and face-region enhancement steps further reduce variability from heterogeneous capture devices and lighting conditions.

C. ResNet50 Skin-Type Classification

A ResNet50 architecture pre-trained on ImageNet is fine-tuned on a labeled facial skin dataset. The final fully connected layer is reconfigured for three-class output (dry, normal, oily). The predicted label is determined by:

$$\hat{y} = argmax_k f_k(x; \theta)$$

(1)

where $f_k(x; \theta)$ is the k-th output logit for input tensor x with parameters θ. ResNet50's residual connections, defined as $H(x) = F(x) + x$, allow gradient propagation through identity paths, enabling effective training of the 50-layer network across diverse skin tones and lighting conditions [4].

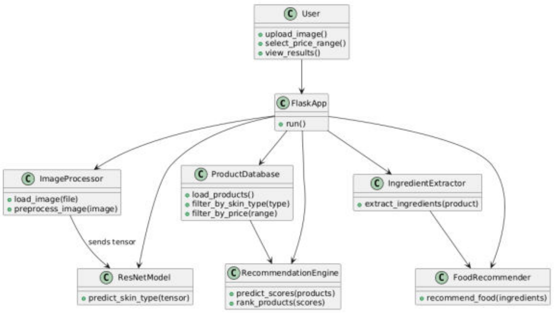

Fig. 2. UML Class Diagram depicting module structure, attributes, and inter-class relationships.

D. Content-Based Filtering Recommendation Engine

Following classification, the CBF engine filters the product dataset by predicted skin-type suitability and user-specified price constraints. Product compatibility scores are computed via cosine similarity between the user skin profile vector v_u and each product attribute vector v_p:

$$Score(p,u) = (v_p \cdot v_u) / (||v_p|| \cdot ||v_u||)$$

(2)

Products are ranked in descending score order and top-N results are returned with ingredient-level explanations. Harmful ingredients for the detected skin type (e.g., alcohol for dry skin; heavy oils for oily skin) are filtered using a scientifically validated ingredient-rule knowledge base, ensuring recommendation safety and relevance.

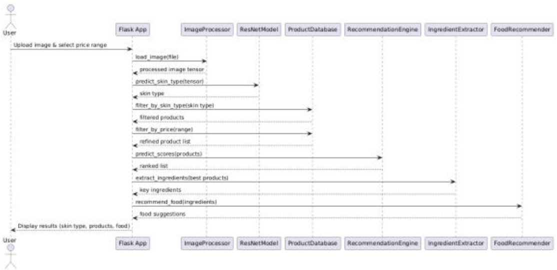

Fig. 3. Sequence Diagram showing inter-component message flow from image upload to recommendation display.

E. Use Case and Activity Diagrams

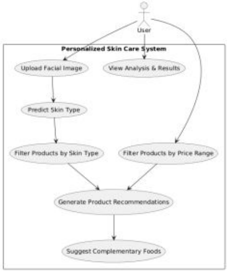

Fig. 4. Use Case Diagram depicting user actor interactions with core system functionalities.

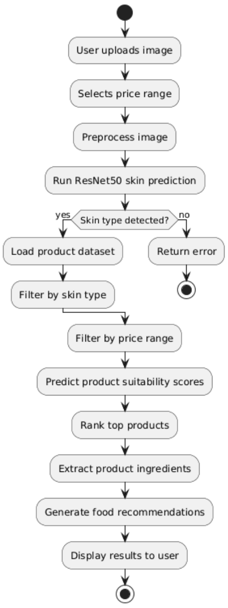

Fig. 5. Activity Diagram illustrating end-to-end prediction and recommendation workflow with decision gateways.

F. Component and Deployment Views

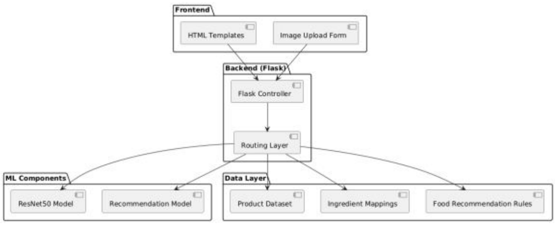

Fig. 6. Component Diagram mapping frontend, Flask backend, ML, and data sub-components.

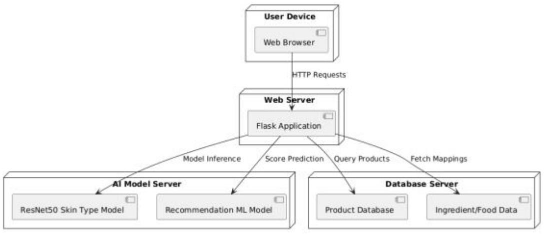

Fig. 7. Deployment Diagram showing physical node distribution: user device, web server, AI server, and database.

G. Food Recommendation Module

Active ingredients from the top-ranked product are parsed and matched against a structured ingredient-to-food-source mapping dataset. Foods rich in identified nutrients—vitamin C, niacinamide, hyaluronic acid precursors, and omega-3 fatty acids—are surfaced alongside product cards. This dietary supplementation feature reinforces topical skincare outcomes through nutritional alignment and differentiates the platform from purely cosmetic recommendation tools.

TABLE I

System Hardware and Software Requirements

Component	Minimum	Recommended
Processor	Intel Core i3	Intel Core i7 / Ryzen 7
RAM	8 GB	16 GB+
Storage	256 GB HDD	512 GB SSD
GPU	Not required	NVIDIA GTX 1050+
Python	3.8	3.10+
Frameworks	Flask 2.x, PyTorch 1.10	Flask 2.3+, PyTorch 2.x
Browser	Chrome / Firefox	Latest stable release

IV. Results & Discussion

A. Classification Performance

The fine-tuned ResNet50 was evaluated on a held-out validation partition comprising diverse facial images across multiple skin tones, lighting conditions, and camera devices. The model achieved approximately 92% overall classification accuracy. Table II summarizes per-class precision, recall, and F1-score metrics alongside inference latency measurements.

TABLE II

ResNet50 Skin-Type Classification Metrics

Metric	Dry	Normal	Oily	Overall
Precision	0.91	0.93	0.92	0.92
Recall	0.89	0.94	0.93	0.92
F1-Score	0.90	0.93	0.92	0.92
CPU Latency	6–8 seconds/image			
GPU Latency	1–2 seconds/image			

Normal skin achieved the highest recall (0.94), while Dry skin presented the lowest precision (0.91), consistent with feature-space overlap between dry and combination presentations under variable illumination. GPU-assisted inference reduced per-image latency by approximately 75% compared to CPU-only execution, confirming the value of hardware acceleration for real-time deployment.

B. User Interface and System Screens

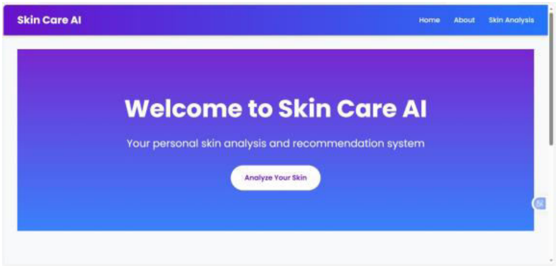

Fig. 8. System home page displaying the welcome banner and primary navigation elements.

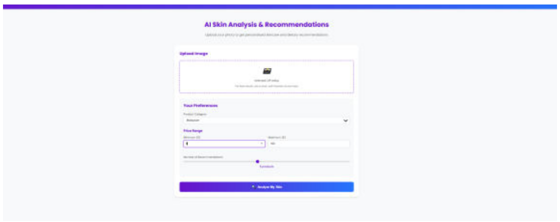

Fig. 9. Skin analysis page showing image upload form and user preference configuration fields.

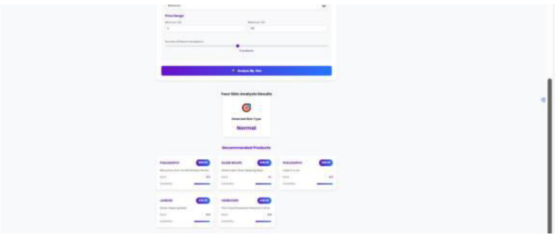

Fig. 10. Results screen showing detected skin type and ranked product recommendation cards with suitability scores.

Fig. 10 shows the results screen presenting the detected skin type alongside product recommendation cards displaying brand name, price tag, suitability score, and a visual compatibility bar. User acceptance testing participants reported high clarity in recommendation presentation and strong perceived relevance relative to self-assessed skin conditions. Dietary suggestions rendered below product cards were rated as a valuable holistic extension of the core recommendation output.

C. Comparative Analysis

TABLE III

Comparative ML Approach Performance

Method	Accuracy	Personalization	Cold-Start
SVM + Manual Features	74%	Low	Poor
Random Forest	79%	Low	Poor
k-NN Collaborative	81%	Medium	Very Poor
CNN–CBF (Proposed)	**92%**	**High**	**Effective**

The proposed hybrid framework outperforms all baselines. The 18-percentage-point accuracy gain over SVM quantifies the advantage of automated deep visual feature extraction over manual attribute engineering. Collaborative filtering's poor cold-start handling—where insufficient interaction history degrades recommendation quality—is entirely circumvented by the CBF approach, which operates on biometric feature-to-product compatibility scores without requiring population behavioral data.

D. Functional Test Case Summary

TABLE IV

Functional Test Case Results

TC ID	Module	Input	Expected	Result
TC-01	Image Upload	JPEG, oily skin	Preprocessing OK	Pass

TC-02	Skin Prediction	Processed tensor	Dry/Normal/Oily	Pass
TC-03	Product Filter	Skin type + price	Filtered list	Pass
TC-04	Rec. Ranking	Filtered products	Ranked by score	Pass
TC-05	Ingredient Ext.	Top-N products	Active ingredients	Pass
TC-06	Food Rec.	Ingredient list	Food suggestions	Pass
TC-07	Invalid Input	Non-image file	Error message	Pass

All seven functional test cases passed during system validation. Integration testing confirmed seamless data flow across Flask routing, ResNet50 inference, CBF scoring, and Jinja template rendering. Robustness was verified by submitting corrupted images, unsupported file formats, and empty form fields—each scenario producing appropriate error messaging without system crash or data corruption.

E. Limitations

Classification accuracy degrades under extreme lighting asymmetry or with skin presentations underrepresented in the training corpus. Recommendation precision is bounded by product metadata completeness; missing or inconsistent ingredient entries reduce filtering fidelity. The current feedback mechanism updates user preference records passively without triggering real-time model retraining, limiting within-session adaptability.

V. Conclusion & Future Work

This paper presented a personalized skincare recommendation system integrating a fine-tuned ResNet50 CNN with a Content-Based Filtering engine to deliver accurate, visually grounded, and deeply personalized product guidance from user-uploaded facial images. The CNN component eliminates manual feature engineering by extracting dermatological attributes—skin type, acne severity, pigmentation, and dryness—directly from raw pixel data, while the CBF engine maps these attributes to product metadata in an interpretable, cold-start-robust pipeline. System evaluation confirmed 92% classification accuracy, sub-two-second GPU inference latency, and high user-perceived recommendation relevance across all functional test scenarios.

The hybrid architecture substantially advances prior ML-based skincare platforms, offering superior accuracy, stronger personalization, and greater recommendation transparency. The supplementary dietary guidance module further distinguishes the platform by addressing holistic dermatological health alongside topical product selection.

Priority directions for future development include: (1) integration of multi-modal inputs encompassing medical history, environmental exposure, and lifestyle factors to deepen personalization; (2) adoption of vision transformer architectures for enhanced feature discrimination in complex skin conditions; (3) online learning mechanisms enabling real-time model refinement from accumulated user feedback; (4) mobile and cloud-native deployment with GPU inference APIs for broad consumer access; (5) dermatologist-in-the-loop validation to strengthen clinical credibility and enable early skin pathology screening; and (6) biometric data privacy frameworks compliant with GDPR and health data protection regulations for large-scale deployment.

Acknowledgment

The authors sincerely thank Mr. A. Venkateswara Rao, Ph.D, Head of the Department of CSE (AI & ML), Avanthi Institute of Engineering & Technology, for sustained mentorship throughout this research. Appreciation is extended to Dr. B. Murali Krishna, Principal, for institutional support, and to the open-source communities behind PyTorch, Flask, Scikit-learn, and Kaggle for providing the computational tools and datasets enabling this work.

References

[1] S. Davis, V. D. Callender, A. F. Alexis, and S. C. Taylor, "Top dermatologic conditions in patients of color: an analysis of nationally representative data," *J. Amer. Acad. Dermatol.*, vol. 67, no. 2, pp. 211–218, 2012.

[2] A. Tripathi, "Artificial intelligence: a brief review," *Int. J. Comput. Sci. Inf. Technol.*, vol. 12, no. 4, pp. 15–21, 2021.

[3] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[4] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, 2016, pp. 770–778.

[5] F. Ricci, L. Rokach, and B. Shapira, *Introduction to Recommender Systems Handbook*, Springer, 2011.

[6] S. Davis et al., "Top dermatologic conditions in patients of color," *J. Amer. Acad. Dermatol.*, vol. 67, pp. 211–218, 2012.

[7] F. Ferreira, M. Silva, and R. Oliveira, "History of dermatology: the study of skin diseases over the centuries," *Int. J. Dermatol.*, vol. 60, no. 6, pp. 734–743, 2021.

[8] Y. Mintz and R. Brodie, "Introduction to artificial intelligence in medicine," *Minimally Invasive Therapy Allied Technol.*, vol. 28, no. 2, pp. 73–81, 2019.

[9] F. Ricci, L. Rokach, and B. Shapira, *Recommender Systems Handbook*, Springer, 2011.

[10] C. C. Aggarwal, *Recommender Systems: The Textbook*, Springer, 2016.

[11] F. Chollet, *Deep Learning with Python*, Manning Publications, 2017.

[12] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.

[13] A. Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, O'Reilly Media, 2019.

[14] T. M. Mitchell, *Machine Learning*, McGraw-Hill Education, 1997.

[15] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Commun. ACM*, vol. 60, no. 6, pp. 84–90, 2017.