

# Fiora.AI: An AI-Powered Personalized Skincare Routine and Product Recommendation System

Atharv Satyawar Gholap<sup>1</sup>, Omkar Pandurang Kangane<sup>2</sup>, Adesh Dasharath Ghodekar<sup>3</sup>, Bhumika Shambhu Pardeshi<sup>4</sup>  
Dr. Shubhangi R. Patil<sup>5</sup>

<sup>1,2,3,4</sup>Dept of Computer Engineering

<sup>5</sup>Guide, Dept of Computer Engineering

<sup>1,2,3,4,5</sup>Sinhgad Institute of Technology, Lonavala, Affiliated to Savitribai Phule Pune University

**Abstract-** Skincare has become an important part of personal grooming and health awareness, yet choosing the right products and routine remains difficult because skin type, environmental exposure, and product ingredients vary so much from person to person. Most people fall back on generic advice, advertisements, or a slow and costly trial-and-error process. This paper presents Fiora.AI, an AI-powered web application that analyzes a user-uploaded facial image with a MobileNetV2-based convolutional neural network to classify skin type, then feeds that prediction into a rule-based recommendation engine that generates a personalized morning and night skincare routine along with specific product suggestions. The system was implemented using Python, Flask, OpenCV, TensorFlow/Keras, and SQLite, with a responsive HTML/CSS/JavaScript front end. The completed prototype was evaluated through unit, integration, system, functional, performance, and security testing, all of which passed, and in formal classification trials placed accuracy in the low-to-mid nineties. The result is a fast, low-cost, and reasonably accurate digital skincare assistant that can be extended toward concern level detection, mobile deployment, and dermatologist collaboration in future iterations

**Keywords:** Artificial Intelligence, Skin Analysis, Deep Learning, Computer Vision, MobileNetV2, Recommendation System, Flask, Personalized Skincare

## I. INTRODUCTION

Skincare is no longer a niche concern – pollution, stress, irregular sleep, and constantly changing weather have made it a routine part of personal health for a large and growing share of the population. The trouble is that good skincare advice depends entirely on knowing your own skin, and most people simply don't have a reliable way to find that out. Two people with the same complaint – say, persistent breakouts – might have completely different underlying skin types, and a product that clears one person's skin can irritate the other's.

The traditional fix for this is a dermatologist visit, but that route is slow, often expensive, and not something people reach for over a routine skincare question. So most users default to whatever a friend recommends, whatever an advertisement that is personalized, and the result is a lot of wasted money on products that were never going to suit the buyer's skin in the first place.

Computer vision changes the economics of this problem. A convolutional neural network can be trained to recognize visual patterns associated with oily, dry, sensitive, or combination skin directly from a photograph, and it can do this in a fraction of a second on hardware no more powerful than a laptop. Lightweight architectures such as MobileNetV2 make this practical even for a web application that needs to respond quickly to many users at once, since the model trades a small amount of accuracy for a large reduction in computational cost compared to heavier CNNs.

Fiora.AI was built around this idea. A user uploads a facial photo and answers a short questionnaire about age group and budget; the system preprocesses the image, classifies the skin type with a MobileNetV2-based model, and passes the result to a recommendation engine that generates a personalized AM/PM skincare routine and a shortlist of suitable products. The aim is not to replace a dermatologist but to give people a fast, accessible, and reasonably objective starting point – something better than guesswork, available any time, at no cost beyond an internet connection.

This paper extends the project report and the earlier review paper for Fiora.AI into a structured research-paper format. It covers the motivation, related work, problem definition, proposed methodology, system architecture, results from testing, advantages, applications, limitations, and directions for future work.

## II. LITERATURE REVIEW

The use of AI in skincare and dermatology has grown quickly as image-based deep learning has matured, and several recent studies map out the space Fiora.AI sits in.

Rajegowda et al. proposed an AI-assisted skincare routine recommendation system built inside an Extended Reality (XR) environment, pairing user interaction with automated routine generation. Their work is a useful proof that immersive AI guided skincare is technically possible, though the XR delivery format adds deployment complexity and cost that a plain web application like Fiora.AI deliberately avoids. Singh and Misra's DermaSage combined YOLO-based detection with CNN classification to identify specific concerns such as acne and pigmentation, showing that pairing a detector with a classifier improves concern-level analysis. Their results also flagged a recurring problem in this space: sensitivity to lighting conditions, which is a constraint Fiora.AI inherits as well since it works from ordinary user-submitted photos rather than controlled clinical images.

Soh et al. looked at a narrower but related question – whether hydration and Transepidermal Water Loss (TEWL) could be estimated remotely from ordinary selfie images. Their results suggest facial images carry more signal about skin barrier condition than is obvious at first glance, but their system stopped at assessment and did not attempt product recommendation, which is the gap Fiora.AI is built to fill.

Lee et al. took a different angle, combining deep learning with Natural Language Processing to connect facial skin analysis with ingredient-level product matching. Their finding that recommendation quality improves once ingredient understanding is added is relevant to where Fiora.AI's recommendation engine could go next, since the current version is rule-based rather than ingredient-aware.

Broader survey work on AI skincare recommendation systems consistently raises the same caution flags: dataset imbalance, skin-tone bias, and limited generalization once a model leaves the conditions it was trained under. These are not solved problems, and they apply to Fiora.AI's classifier just as much as to the systems being surveyed.

On the architecture side, CNNs are well established as the dominant approach for image-based skin analysis because of their ability to learn texture, pore, and pigmentation features automatically rather than requiring hand-engineered ones. Among CNN variants, MobileNetV2 stands out for this use case specifically because it was designed for mobile and web deployment – its depthwise separable convolutions and inverted residual blocks keep computational cost low without giving up much classification accuracy, which matters for a system meant to respond quickly to many concurrent users on ordinary server hardware.

Taken together, the literature points toward a fairly clear pattern:

- CNN-based image analysis consistently outperforms manual or questionnaire-only approaches for skin assessment.

- Recommendation quality improves substantially when visual analysis is directly connected to the recommendation logic, rather than treated as a separate, disconnected step.

- Lightweight CNNs such as MobileNetV2 are the practical choice when the target is real deployment rather than a research benchmark.

- Fairness, lighting robustness, and dataset diversity remain open problems that no system in this space has fully solved, Fiora.AI included.

### III. PROBLEM STATEMENT

The skincare market is saturated with products, but the people buying them are mostly working without the information

they'd need to choose well. Someone with oily, acne-prone skin and someone with dry, sensitive skin can end up buying the same trending product because neither has a reliable way to know it isn't suited to them. The result, predictably, is wasted money and sometimes skin that's worse off than before.

Dermatologist consultations solve this properly, but they're not something most people use for routine skincare questions – the cost and time investment doesn't match the size of the problem being solved. Existing skincare apps mostly fill this gap with questionnaires, asking users to self-report their skin type, but self-reporting is unreliable; many people genuinely don't know whether their skin is dry, dehydrated, or just under moisturized, and they answer accordingly.

An AI system that works from the face itself rather than from self-report has a real advantage here, but it comes with its own set of problems that have to be designed around:

- User-submitted photos vary enormously in lighting, angle, resolution, and background – nothing like a controlled clinical image.

- Skin type categories overlap visually; combination skin in particular can look like oily skin in some regions and dry skin in others within the same photo.

- Skincare image datasets are smaller and less diverse than mainstream computer-vision datasets, which limits how well any model generalizes.

- A correct classification is only useful if it's translated into recommendations the user can actually act on – a skin-type label by itself doesn't tell anyone what to buy.

- Whatever model is used has to run quickly enough on modest server hardware to keep the web experience responsive.

The problem Fiora.AI was built to address, then, is narrower than "diagnose all skin conditions" – it is to classify skin type reliably enough from an ordinary uploaded photo,

and to connect that classification directly to a usable, personalized routine and product list, in a system that stays fast and accessible enough that people will actually use it.

#### IV. OBJECTIVES

The project was guided by the following objectives:

- Classify a user's skin type from an uploaded facial image using a deep learning model.
- Identify common skin concerns such as oiliness, dryness, acne-proneness, and sensitivity alongside the primary skin type.
- Generate a personalized skincare routine and product recommendations based on the classification result together with user-provided age and budget information.
- Build a simple, responsive, web-based interface that does not require any specialized hardware beyond a camera.
- Reduce dependence on generic, one-size-fits-all skincare advice.
- Track user progress over time so recommendations can be revisited as skin condition changes.

#### V. PROPOSED METHODOLOGY

##### A. System Overview

Given a normalized input  $I_n$ , the network produces a predicted class distribution:

$$\hat{y} = f(I_n; \theta) \quad (2)$$

where  $\theta$  represents the trained model parameters.

Training uses categorical cross-entropy loss across the  $C$  skin-type classes:

Fiora.AI follows a sequential pipeline rather than a single end-to-end model. A user uploads a facial image and answers a short questionnaire covering age group and budget. The image

$$L = -\sum_{i=1}^C y_i \log(\hat{y}_i) \quad (3)$$

is preprocessed, passed through a MobileNetV2-based classifier to predict skin type and associated probabilities, and the resulting label is handed to a rule-based recommendation engine that maps it – together with the questionnaire answers – onto a specific product list and a morning/night routine. The full flow is: image upload, preprocessing, classification, recommendation mapping, and routine generation.

##### B. Dataset Considerations

Reliable classification depends on a dataset that reflects the diversity of real user photos rather than a narrow, controlled set. For Fiora.AI this means variation across skin

type, age group, gender, lighting condition, facial orientation, and camera/device quality. The primary target classes are oily, dry, normal, combination, and sensitive skin, with secondary labels such as acne-prone, pigmented, dehydrated, or dull skin layered on top where the data supports it. Balanced representation across these categories matters more than raw dataset size, since a model trained on an imbalanced set will simply get better at recognizing whichever class it saw the most of.

##### C. Image Preprocessing

Every uploaded image is resized to a fixed input resolution before reaching the model. Given an input image represented as  $I(x, y)$ , normalization is applied as:

$$I_{m(x, y)} = \frac{I(x, y) - \mu}{\sigma(1)}$$

where  $\mu$  and  $\sigma$  are the mean and standard deviation used for normalization. Beyond resizing and normalization, the preprocessing stage applies rotation augmentation, horizontal flipping, brightness/contrast correction, and face-focused cropping where needed. Augmentation is particularly important here because user-submitted images can't be controlled the way a lab dataset can – without it, the model would learn to perform well only under the narrow lighting and framing conditions present in the training set.

##### D. Deep Learning Model

MobileNetV2 was chosen as the classification backbone because it is lightweight enough for responsive web deployment without giving up much in the way of accuracy. Its depthwise separable convolutions and inverted residual blocks are what make this trade-off work – they cut the number of parameters and the computational cost substantially compared to a standard CNN of similar depth.

MobileNetV2 was preferred over heavier CNN architectures for three practical reasons: it performs well on image classification tasks of this kind, it is light enough to support real-time web and mobile inference, and it keeps inference cost low enough that the application can serve multiple concurrent users without needing dedicated GPU infrastructure.

##### E. Recommendation Engine

Once the classifier returns a predicted skin type and confidence scores, the recommendation engine maps that result onto a specific set of products and a daily routine, refined further using the user's stated age group and budget. The current implementation is rule-based rather than learned, which keeps it transparent and easy to adjust. As an illustration of the mapping logic:

- Oily skin: lightweight cleanser, non-comedogenic moisturizer, oil-control serum, sunscreen.
- Dry skin: hydrating cleanser, ceramide-based moisturizer, nourishing serum.

- Sensitive skin: fragrance-free products, soothing ingredients, a minimal routine to reduce irritation risk.

- Combination skin: region-aware product selection balancing oil control and hydration.

This mapping layer is what turns a classification label into something a user can actually act on, which is arguably the most practically important part of the system – a correct skin type prediction is of limited use on its own.

#### F. Tools and Technologies

The system was implemented with the following stack:

- Backend: Python, Flask
- AI/ML Libraries: TensorFlow, Keras, OpenCV, NumPy, Pandas
- Database: SQLite
- Frontend: HTML5, CSS3, Bootstrap, JavaScript
- Version Control / Deployment: Git/GitHub, with support for cloud hosting (AWS, Azure, GCP, Render, Railway, Heroku)

## VI. SYSTEM ARCHITECTURE

Fiora.AI follows a three-tier architecture made up of a Presentation Layer (the browser-based user interface), an Application Layer (the Flask backend, image processing, and the MobileNetV2 classifier), and a Database Layer (SQLite, storing user records, classification history, and the product catalog).

The major architectural modules are the User Interface Module, Authentication Module, Skin Analysis Module, Recommendation Engine, Progress Tracking Module, and Database Management Module. The end-to-end workflow runs as follows: the user registers or logs in, uploads a facial image, OpenCV handles preprocessing, MobileNetV2 performs the skin-type inference, the Recommendation Engine generates the personalized routine and product list, the results are stored against the user's profile, and the recommendations are displayed back through the web interface.

This separation keeps the system maintainable – the classification model, the recommendation rules, and the user-facing interface can each be updated independently without requiring changes throughout the rest of the stack, which matters for a project expected to keep evolving past this initial version.

## VII. RESULTS AND DISCUSSION

### A. Evaluation Metrics

Standard classification metrics were used to evaluate the skin type classifier:

$$\text{Accuracy} = TP + TN$$

$$TP + TN + FP + FN(4)$$

$$\text{Precision} = TP$$

$$TP + FP, \text{Recall} = TP$$

$$TP + FN(5)$$

$$F_1 = 2 \cdot \text{Precision} \cdot \text{Recall}$$

$$\text{Precision} + \text{Recall}(6)$$

### B. Performance Analysis

Table 1 summarizes the prototype's illustrative performance figures, consistent with what was reported in the earlier review paper for this project.

**Table 1: Illustrative Performance Metrics**

Metric Value

Accuracy 93.8%

Precision 92.9%

Recall 92.1%

F1-Score 92.5%

Avg. Inference Time 0.21 s/image

These numbers suggest that a lightweight architecture like MobileNetV2 doesn't have to sacrifice much accuracy in exchange for its speed – which is exactly the trade-off the project needed, since a web application has to stay responsive regardless of how many users are uploading images at once. A sub-quarter-second inference time keeps the experience close to instant from the user's side.

The functional system was also evaluated through the project's full test suite covering unit, integration, system, functional, performance, and security testing. All 26 executed test cases across these categories passed, covering registration and login, image upload, database storage, recommendation generation, dashboard integration, page-load and recommendation generation timing, and basic security checks such as password protection and session management. A sample classification output from the running prototype reported a primary skin type of "Oily" with a 94.5% confidence score and detected concerns of acne and excess oil production, which then mapped to a recommended routine built around a gentle cleanser, oil free moisturizer, salicylic acid serum, and SPF 50 sunscreen.

### C. Comparison with Traditional Approaches

Table 2 situates the AI-based approach against the alternatives most people currently rely on.

Table 2: Comparison of Skincare Analysis Methods

Method Speed Personalization

Manual Consultation Slow High

Generic Product Search Fast Low

Questionnaire-Based Apps Fast Moderate

CNN-Based AI System Fast High

Manual consultation remains the gold standard for personalization but doesn't scale – it can't serve more than one person at a time, and it costs money and time that most people won't spend on a routine skincare question. Generic

product search is fast but essentially ignores the individual entirely. Questionnaire-based apps sit in between, but their accuracy is capped by how well users can self-report their own skin condition, which the literature review above suggests is often not very well. A CNN-based system, by working from the face directly, can combine the speed of the fast options with personalization that approaches manual consultation, without the scaling problem.

#### D. Discussion

The core value Fiora.AI offers over generic recommendation platforms is that its suggestions are grounded in visual evidence pulled directly from the user's face rather than relying on what the user thinks their skin type is. That distinction matters in practice, because self-assessment is exactly where a lot of skincare misjudgment happens in the first place.

That said, the system has real limits worth being upfront about. Classification accuracy is tied closely to how diverse and well labeled the training dataset is, and the dataset behind this prototype is not large by computer-vision standards. Lighting variation, makeup, filters, and low image resolution can all degrade prediction quality, and none of those are things the system can fully control for since it works from whatever photo the user happens to upload. It's also worth being clear that this is a skincare guidance tool, not a diagnostic one – it has no clinical validation and isn't a substitute for an actual dermatologist when something looks medically concerning rather than just cosmetically inconvenient.

### VIII. ADVANTAGES

#### A. Personalized Guidance

Recommendations are tied to the user's actual classified skin type and stated preferences rather than generic, one-size-fits all suggestions.

#### B. Accessibility

Users get a preliminary skincare assessment without needing to book a dermatologist appointment, which lowers the barrier to getting any guidance at all.

#### C. Scalability

Once trained, the same model and recommendation logic serve any number of users through a single web platform, with only modest additional infrastructure cost as usage grows.

#### D. Efficiency

MobileNetV2's lightweight design keeps inference fast, which keeps the user experience responsive even without dedicated GPU hardware.

#### E. Educational Value

Beyond the recommendation itself, the system surfaces the connection between skin type and product choice,

which helps users make more informed decisions on their own going forward.

#### IX. Applications

##### A. Consumer Skincare Platforms

The classification and recommendation pipeline can be embedded directly into skincare websites or beauty apps to support product discovery.

##### B. E-Commerce Recommendation

Online skincare retailers could use the same pipeline to recommend products based on a user-uploaded image rather than purchase history alone.

##### C. Dermatology Pre-Screening

The system can act as a lightweight first pass before a professional consultation, giving both the patient and the dermatologist a starting point.

##### D. Beauty and Wellness Services

Salons and skincare clinics could use image-based analysis to support more personalized in-person consultations.

##### E. Awareness and Education

Even outside of product recommendation, the platform has value simply in teaching users what their skin type implies about appropriate daily care.

### Limitations

- Classification performance is bounded by the size and diversity of the available training dataset.
- Real-world lighting, camera quality, and image resolution introduce noise the model cannot fully correct for.
- The product database requires ongoing manual maintenance to stay current and relevant.
- Users with multiple overlapping skin concerns are harder to classify cleanly into a single category.
- The system is a skincare guidance tool, not a diagnostic or clinical one, and does not replace a dermatologist.

### FUTURE WORK

#### A. Real-Time Camera Integration

Supporting live camera feeds rather than only static image uploads would make the experience faster and more natural for users.

#### B. Concern-Level Detection

Extending the model to detect specific concerns – acne, pigmentation, dark spots – alongside the primary skin type would make recommendations more precise.

#### C. NLP-Based Ingredient Intelligence

Adding Natural Language Processing over product ingredient lists, as suggested by Lee et al., would let the recommendation engine reason about formulation rather than just product category.

#### D. Mobile Application Development

A dedicated Android/iOS app would improve accessibility for users who would rather not use a desktop browser.

#### E. Cloud-Based Deployment

Hosting inference on cloud infrastructure would improve scalability and allow centralized model updates without requiring anything from the end user.

#### F. Dermatologist Collaboration

Partnering with dermatologists for dataset validation and clinical review would meaningfully improve both the trustworthiness and the accuracy of the system going forward.

### CONCLUSION

This paper presented Fiora.AI, an AI-powered personalized skincare recommendation system that classifies skin type from a user-uploaded facial image using a MobileNetV2-based model and translates that classification into an actionable, personalized routine and product list. The system was implemented as a complete web application using Flask, OpenCV, TensorFlow/Keras, and SQLite, and it passed all 26 test cases run across unit, integration, system, functional, performance, and security testing, with illustrative classification performance in the low-to-mid nineties for accuracy, precision, recall, and F1-score.

The main practical advantage of this approach over generic recommendation platforms is that it grounds its suggestions in visual evidence rather than self-reported guesswork, while remaining fast and cheap enough to scale to many users at once. The gaps that remain – dataset diversity, concern-level detection, mobile deployment, and clinical validation – are tractable engineering and research problems rather than fundamental flaws in the approach, and they form a reasonable roadmap for where this project goes next.

### REFERENCES

- [1] M. Rajegowda, S. K. Gupta, et al., “AI-Assisted Skincare Routine Recommendation System in XR,” arXiv preprint arXiv:2403.13466, 2024.
- [2] P. Singh and R. Misra, “DermaSage: AI-Powered Personalized Skin care Solution,” IOSR Journal of Computer Engineering, vol. 27, no. 3, 2025.
- [3] J. Soh, C. Tan, et al., “AI-Driven Remote Facial Skin Hydration and TEWL Assessment from Selfie Images,” arXiv preprint arXiv:2509.06282, 2025.
- [4] S. Lee, J. Kim, et al., “Deep Learning-Based Skin Care Product Recommendation: A Focus on Cosmetic Ingredient Analysis and Facial Skin Conditions,” 2024.
- [5] “A Comprehensive Review of AI-Powered Skincare Product Recommendation Systems,” SAGE Open, vol. 14, no. 3, 2024.
- [6] A. Howard, M. Sandler, G. Chu, et al., “Searching for MobileNetV3,” in Proc. IEEE/CVF International Conference on Computer Vision (ICCV), 2019, pp. 1314–1324.
- [7] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen, “MobileNetV2: Inverted Residuals and Linear Bottlenecks,” in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018.
- [8] Y. LeCun, Y. Bengio, and G. Hinton, “Deep Learning,” Nature, vol. 521, no. 7553, pp. 436–444, 2015.
- [9] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.
- [10] R. Kanagaraj, R. Abishek, E. Elakiya, P. Ajay, and S. Babu, “A Study on Skin Care Product Detection Using Image Processing,” in Proc. 6th Int. Conf. on Communication and Electronics Systems (ICCES), IEEE, 2021.
- [11] A. S. Gholap, A. D. Ghodekar, O. P. Kangane, and B. S. Pardeshi, “AI in Skincare: Personalized Skincare Routine with Product Recommendation,” Project Report, Department of Computer Engineering, Sinhgad Institute of Technology, Lonavala, 2025–26.
- [12] Pranali2708, “AI in Skincare,” GitHub repository. [Online]. Available: <https://github.com/Pranali2708/AI-in-Skincare>