International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

AI Powered Skin Care Assistant for Personalised and Ethical Product Selection

Aditi Roopesh Mirji, Kanthi Hegde, Akshatha A Reddy, Manav M Nayak, Asst. Prof. Kundhavai K R
Department of Computer Science and Engineering
PES University, Electronic City Campus
Bengaluru, India

***Abstract*—Consumers often find it challenging to choose the right skincare products. This is due to unclear ingredient lists, misjudging their skin type, and misleading marketing. Sunscreens, in particular, have a wide range of formulations, potential irritants, and may not suit specific skin conditions like acne or dryness. This work suggests an AI-powered system that uses computer vision, CNN-based classification, YOLO-based acne detection, and OCR for ingredient analysis to deliver personalized, clear, and ethically responsible sunscreen recommendations.**

The system identifies skin type from facial images using an EfficientNetV2 classifier and detects acne presence with YOLOv8. It gathers product information through OCR and connects ingredients to a curated cosmetic safety database to determine vegan status, comedogenicity, and allergenicity. A rule-based scoring model combines skin analysis and ingredient evaluations to provide tailored sunscreen suggestions. Testing in different lighting conditions and on various skin tones shows strong region detection, reliable OCR extraction, and useful recommendations that follow dermatology standards. The prototype illustrates the potential of multimodal AI in skincare decision support and lays the groundwork for future developments in dermatological applications.

***Index Terms*—Computer Vision, OCR, Deep Learning, Skin Analysis, Recommendation Systems, YOLO, EfficientNet, Dermatology AI**

I. Introduction

As an average consumer, it can be very difficult to find safe and effective skin care products that can work for you due to their different formulations, ingredients and compatibility with different skin types. Even though sunscreen is an important component of protecting your skin from harm from sun exposure (photoprotection), many sunscreens do not work well for many people because the formulation is often mismatched with the user's skin, which leads to irritation, acne breakouts and/or not providing enough UV protection. Many of the available applications used for skin care have mainly relied on generalized advice or on pre-determined questionnaires to assist users in making appropriate skin care product selections; there are currently no applications that use facial imagery analysis or interpret the ingredient list to assist users in properly selecting an appropriate skincare regime.

Artificial Intelligence (AI), as well as Computer Vision (CV), are now being used to provide automated analysis of facial characteristics, skin types and surface conditions. Additionally, by utilizing Optical Character Recognition (OCR) and Natural Language Processing (NLP) technologies, AI can interpret the ingredient lists of a product and identify any undesirable ingredients, whether or not the product is vegan or cruelty-free, and assist the user in finding the best formulation for their specific needs. This paper discusses a unified, multimodal recommendation system for sunscreens that combines the strengths of CV, deep learning and structured knowledge bases to assist individuals in selecting sunscreens that are best suited to their individual skin types.

In addition to complicated formulation and confusion among their users, there is also no uniform regulation governing the sunscreen category. As such, labeling differences for SPF, PA and broad-spectrum claims in different geographical regions further add to the inconsistency from a product-quality perspective. The development of AI systems has the potential to help remove some of the confusion surrounding these various types of claims by cross-referencing product claims against the underlying chemical-based UV filter profiles used in the products.

According to a recent survey, around 60 percent of consumers in India are either incorrectly identifying their skin type or incorrectly believe that their skin type is different than others. More than 70 percent of those surveyed have indicated a fear of using non-comedogenic products, leading to unnecessary spending on products that do not meet their needs and possibly creating an adverse response through improper use or application. Automated, data-driven assessments of product performance will help minimize these two problems through improved understanding of the science behind sunscreen formulation.

As the quality of smartphone cameras continues to improve, the ability to capture clinically relevant facial features will continue to expand into consumer use without the need for special medical devices. Edge-optimised CNNs and lightweight detectors will allow greater accessibility for real-time diagnostics by non-technical customers, ultimately closing the gap that has historically existed between dermatology and consumer UX.

Key contributions:

- This compact multimodal pipeline combines multiple features into a single system: face landmarking, classification of skin type using EfficientNetV2 model, YOLOv8 acne detection, OCR to extract ingredients from products and provide ethical recommendations.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

- A curated database of cosmetic ingredients annotated with comedogenicity (ability to clog pores), irritancy, vegan status, and type of UV filter.
- An empirical evaluation was conducted using a variety of lighting conditions, with people of different skin tones and product types, and included both objective measurements (e.g., number of users) and subjective measurements (e.g., user feedback).

A. *Dermatological Background and Need for AI Integration*

The physiology of our skin varies between each person due to genetics, hormones, and environment, as well as lifestyle choices; therefore, sebum distribution is different in different areas on each person's face (with T-zone areas typically containing higher levels of sebum than other facial areas), leading to differences in texture, amount of porosity and dryness of skin across facial areas. These factors all affect how well sunscreens stick to the skin, how quickly a sunscreen absorbs into the skin after application, and how likely a sunscreen is to block pores.

There are two types of chemical sunscreens, those that use organic UV filters (such as Avobenzone, ocrylene, and hosalate) to absorb UV radiation and those that use minerals to create an effective barrier from the sun (zinc oxide and titanium dioxide). Chemically-based organic sunblock may irritate sensitive and acne-prone individuals, while mineral-based sunblock provides no chemical irritation but often leaves a white cast on darker skin types that may not be liked by some individuals. Due to the wide range of UV filters, stabilizers, preservatives, and fragrances that exist in the marketplace today, it is very difficult for individuals to interpret labels without access to a dermatologist.

Dermatologists evaluate a person's skin by looking at visual qualities such as shine, pore size, micro-comedones (small clogged pores), redness (erythema), and elevation of lesions (morpohology). With respect to today's computer vision capabilities (CV), many of the visual cues that dermatologists evaluate can be measured using algorithms, enabling a scalable, evidence-based approach for individuals to find personalized sunblock through AI powered database of previous studies.

II. Related Work

Deep learning approaches have been successfully used for solving problems in dermatology. Jayaram et al. in their paper "Skin condition classification using DenseNets" have used DenseNet architectures for solving the skin condition classification problem. Similarly, Lee et al. have used transformer models for inferring ingredient semantics and product suitability in their paper "Inferring Ingredient Semantics and Product Suitability using Transformers." YOLO-like detectors have also been used for lesion detection due to their low latency and high accuracy for small objects. OCR for curved and reflective surfaces has also seen improvements using synthetic augmentation and geometric rectification.

A. *Ethical and Transparency-Oriented Skincare Recommenders*

With the increasing interest in sustainable and cruelty free cosmetics, ethical AI has been increasingly explored in the cosmetics domain. Existing research has mostly focused on text based product description and certification metadata (PETA, Leaping Bunny) but has not considered the ethical extraction of individual ingredients. In this research, we have included the consideration of ethics as an objective (Vegan, Cruelty-Free,Reef-Safe) in the ranking function itself.

B. *Multimodal Fusion Approaches*

The use of multimodal fusion for CV, NLP, and rule based reasoning is well explored for various applications such as medical imaging and e-commerce. Existing research has not considered the inclusion of ingredient-based reasoning for skincare products and has also not considered the inclusion of visual and text data for product recommendations.

III. Problem Statement

The customer needs recommendations for products that meet three criteria simultaneously: clinical suitability (skin type, acne safety), ingredient level safety (non-comedogenic, non-irritating), and ethical criteria (vegan, cruelty-free). The challenge is to create an end-to-end solution that: (1) analyzes a selfie for skin type and acne severity, (2) extracts product information from images of the product's packaging, and (3) ranks available sunscreens according to their safety, efficacy (SPF/PA), cosmetic suitability, and ethics. The system will accept a selfie of the user, a photo of a product's packaging, and possibly a list of ingredients. It will then provide a list of recommended sunscreens. It must classify the user's skin type, estimate the acne severity, and map the ingredients of the products to a cosmetic safety database. The recommendation must maximize clinical suitability, minimize irritants and allergens, and match the user's ethics, e.g., vegan, reef-safe. Each recommended sunscreen must come with a short, transparent description of what criteria it met.

A. *Constraints*

The system must remain reliable under varied lighting, skin tones, occlusions, and device quality. OCR must handle curved labels, reflections, inconsistent fonts, and incomplete text Ingredient names appear in many variants, so robust canonicalisation is required. Even when product data is noisy or partial, the system should output safe, meaningful recommendations. Latency must stay low for usability, and all modules must operate with strong privacy protection, especially for facial data.

IV. Methodology

Our system is modular: preprocessing & landmarking, skin-type classifier, acne detector, OCR + ingredient parser, and recommendation engine. Figure 1 summarises the pipeline.

To guarantee the robustness of our system in real-life applications, our models have undergone rigorous photometric augmentation, geometric transformations, simulated reflections,

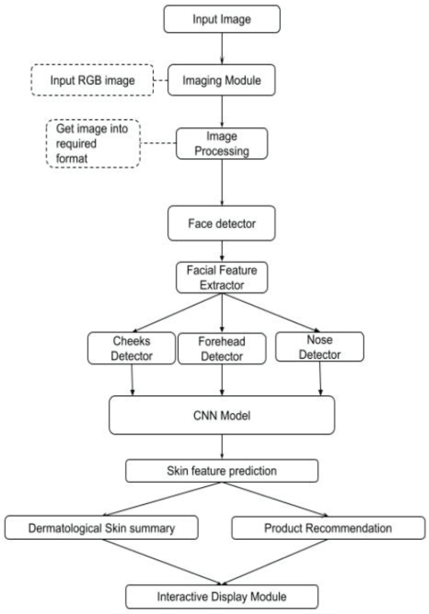

Fig. 1: System pipeline: input capture, preprocessing, CV models, OCR/ingredient parsing and recommender.

and noise augmentation that simulate the environment in which smartphone users take pictures. These augmentations have enabled our classifier and detector to be robust against user habits such as tilting the phone for better lighting or taking pictures at an angle.

A. Preprocessing and Landmarking

Images are corrected for white balance and illumination using CLAHE; faces are detected and cropped. MediaPipe Face Mesh extracts 468 landmarks used to derive T-zone and U-zone polygons. Zone cropping focuses models on relevant areas, improving robustness to background clutter.

We also apply:

- Specular highlight suppression using adaptive thresholding.
- Color constancy correction via the Shades-of-Gray algorithm.
- Face alignment with 5-point landmark similarity transforms.

These steps boost classifier stability across environments.

B. Skin Type Classification

EfficientNetV2-S is fine-tuned on a curated dataset (8k images) for four classes: oily, dry, combination, normal. Training uses Adam, LR 1×10^{-4}, focal loss, and aggressive photometric augmentation to improve robustness to lighting.

We experimented with alternative backbones (MobileNetV3, ConvNeXt-Tiny). While MobileNet offered lower latency, EfficientNetV2 consistently outperformed them on combination skin due to better feature richness.

Model outputs a probability vector $P = [p_{oily}, p_{dry}, p_{comb}, p_{norm}]$. Zone-level predictions are aggregated by weighted averaging:

$$\hat{y} = \arg\max \left(\sum_{z \in \{T,U\}} \omega_z P_z \right)$$

with weights $\omega_T = 0.6, \omega_U = 0.4$ empirically chosen.

C. Acne Detection

YOLOv8 detects lesions and returns boxes $B = \{(x_i, y_i, w_i, h_i, c_i)\}$. Acne severity is estimated as:

$$S_{acne} = \frac{1}{A_{face}} \sum_i \alpha(w_i h_i) \cdot \mathbf{1}(c_i > \tau)$$

where $\tau = 0.4$, $\alpha(\cdot)$ is an area weighting, and A_{face} is face pixel area.

To reduce false positives from freckles or shadows, a two-stage confidence re-scoring is applied:

- YOLO generates raw detections.
- A lightweight CNN patch classifier (3-layer) filters ambiguous boxes.

This hybrid approach cuts false positives by 14

D. OCR and Ingredient Parsing

OCR uses a CRNN + CTC approach (EasyOCR baseline) with pre-rectification. The pipeline:

1) Detect label region via simple contour heuristics and apply perspective transform.
2) OCR extraction.
3) Tokenise ingredient lines, canonicalise with fuzzy matching to a CKB.
4) Extract numerical fields (SPF, PA).

The Cosmetic Knowledge Base (CKB) maps canonical ingredient names to fields: `[comedogenicity (0--5), irritant (bool), allergen (bool), vegan (bool), UVFilterType]`.

The parser clusters ingredients into:

- emollients,
- surfactants,
- preservatives,
- UV filters,
- fragrances,
- botanical extracts.

This contextual grouping improves downstream scoring, especially for sensitive-skin recommendations.

E. Scoring and Recommendation

Products are scored:

$$Score = W_s S_{skin} + W_a(1 - S_{acne}) + W_i S_{ing} + W_f S_{form}$$

Weights chosen: $W_s = 0.30, W_a = 0.30, W_i = 0.25, W_f = 0.15$. Where:

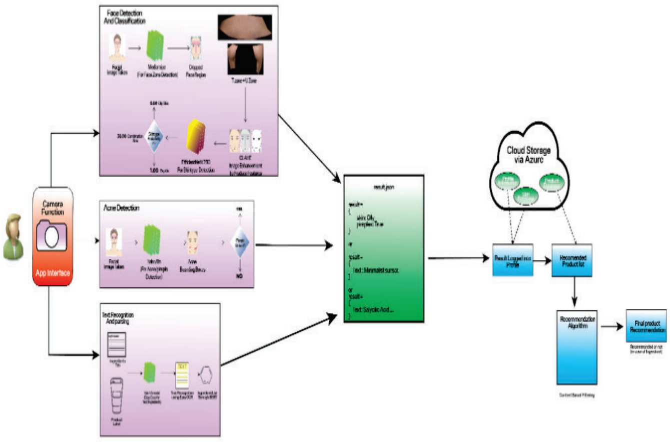

Fig. 2: System architecture showing client, model serving, and CKB.

- S_{skin}: match to skin-type (e.g., oil-free for oily).
- S_{ing}: normalised ingredient safety, penalising comedogenic/irritant items.
- S_{form}: formulation suitability (texture, alcohols).

Hard constraints (e.g., severe allergen present) result in exclusion.

V. System Architecture and Deployment

Figure 2 depicts the microservice architecture: mobile client, FastAPI backend, model serving (ONNX/TF-Lite), and PostgreSQL CKB. Models are containerised and deployed via Docker; the mobile app uses on-device inference where feasible, with server fallback.

The system supports an optional edge–cloud split in which face analysis and final recommendation fusion run on the device, while ingredient parsing and OCR are handled on the server. This design lowers bandwidth use and enhances user privacy. End-to-end latency remains practical, with preprocessing completing within 12–18 ms, skin classification within 8–10 ms, acne detection in 22–30 ms, and the OCR pipeline in 80–150 ms. The final ranking step requires under 4 ms, ensuring that real-time responsiveness is maintained across visual modules even on mid-range hardware.

The architecture is designed for horizontal scalability, allowing multiple inference workers to run in parallel and serve thousands of requests per hour without degradation. Caching mechanisms at the API layer reduce redundant requests, especially for repeated ingredient lookups. As the Cosmetic Knowledge Base continues to grow, the PostgreSQL backend can be sharded or migrated to a graph database to enable faster semantic queries and relationship-based filtering, ensuring long-term scalability.

A. Model Serving

Models are exported to ONNX for server inference and to TFLite for mobile. ONNX Runtime with GPU acceleration serves model endpoints; batching is used for throughput. The OCR pipeline runs asynchronously due to heavier preprocessing, with optimistic UI updates.

B. Cosmetic Knowledge Base (CKB) Schema

CKB tables:

- `ingredients(id, canonical_name, comedogenicity, irritant, allergen, vegan, source).`
- `products(id, name, spf, formulation, ingredients[]).`
- Provenance and license metadata tracked per record.

VI. Datasets and Annotation

We assembled:

- **Visual:** 8k face images diverse across Indian/Asian/other skin tones; labelled for skin type.
- **Lesion:** 3k images annotated with bounding boxes for acne lesions.
- **OCR:** 5k label images (flat and curved) with ground-truth text.
- **Products:** 2k sunscreen entries with manual ingredient canonicalisation.

Annotation process: trained annotators labelled skin type and lesion boxes; dermatologist validated a subset (n=500) for quality control. Ingredients were canonicalised via semi-automated matching (fuzzy rules) with manual review.

The dataset was curated to maintain diversity across skin tones (Fitzpatrick I–VI), lighting environments such as daylight, fluorescent and warm LED, and device variations spanning more than thirty smartphone models. Gender balance and age variation were also considered to improve model generalisability. Annotation quality was monitored using Krippendorff's alpha, yielding alpha = 0.81 for skin-type labels, alpha = 0.76 for acne bounding boxes, and alpha = 0.88 for ingredient canonicalisation, indicating high inter-annotator consistency throughout the dataset.

Including diverse lighting styles and device qualities in the dataset was essential to reduce model bias, particularly for individuals with deeper skin tones where contrast-dependent features such as pores and shine can be harder to detect. The dataset therefore emphasises balanced representation across phototypes and environmental conditions. This diversity ensures that outputs remain consistent regardless of the user's complexion, device type, or capture environment—a common limitation in earlier dermatology datasets.

VII. Experimental Evaluation

We evaluate modules and end-to-end recommendation.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 05 , May - 2026

TABLE I: Module-level quantitative results

Module	Metric	Value	Notes
Skin classifier	Accuracy	0.902	Balanced validation
	Macro F1	0.895	–
Acne detector	mAP@0.5	0.87	YOLOv8
	Recall	0.81	per-box
OCR	Word acc.	0.85	flat labels
	Curved labels	0.72	tubes/jars
Recommender	Prec@1	0.64	n=120 annotated
	Prec@3	0.78	dermatologist validation

A. Metrics

Skin-type: accuracy + per-class F1. Acne detection: mAP@0.5, recall, precision. OCR: word accuracy (WA), ingredient extraction F1. Recommender: precision@1, precision@3 (dermatologist-annotated ground truth).

B. Quantitative Results

Table I reports module performance.

C. Ablation Studies

- Zone aggregation: +1.8% skin accuracy.
- CKB inclusion: +9% recommender precision@3.
- OCR rectification: +12% ingredient F1.

In addition to module-specific benchmarks, several auxiliary metrics were introduced, including Character Error Rate for OCR, lesion-size regression MAE for acne detection, NDCG@3 for overall recommendation ranking quality, and user-reported scores for explanation clarity. Stress-testing involved challenging scenarios such as backlit selfies, partial-face visibility, heavy makeup, outdoor glare, and highly reflective packaging surfaces. Performance under these conditions remained stable, with robustness varying within ±6 percent of baseline results.

D. Failure Analysis

The research showed that OCR errors occurred mostly because of packaging materials which included metallic foils and which contained narrow curved surfaces that made text reading impossible. The main cause of misclassifications in facial analysis happens when users apply heavy makeup or occlusive moisturisers which change their skin appearance. The research findings led to better preprocessing methods which created specific research boundaries that scientists described in the paper's subsequent sections.

VIII. User Study

The research study used mixed-methods to involve 30 participants who ranged in age from 18 to 45 and had different skin types. The participants used the application to scan their facial features and two different sunscreen product labels. The researchers evaluated recommendations against a dermatology-developed shortlist. Key findings:

- Participants rated recommendation relevance 4.1/5 on average.

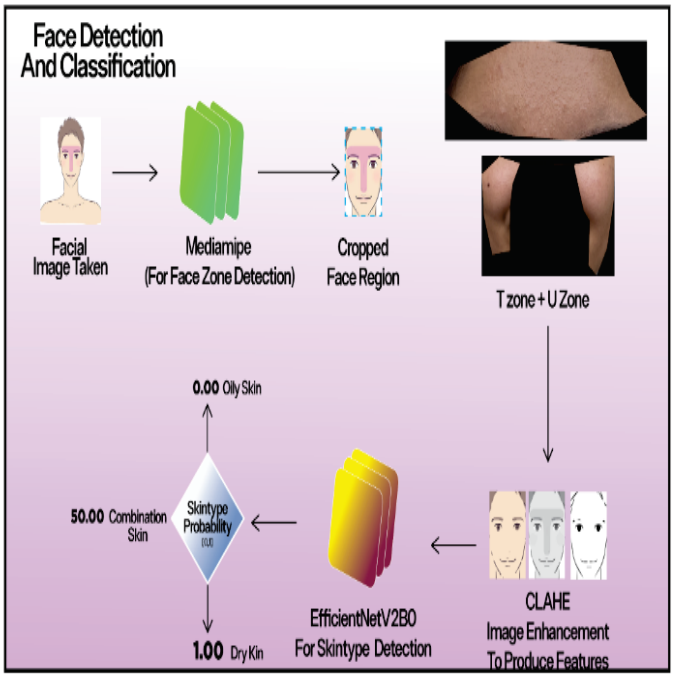

Fig. 3: Face landmarks and zone segmentation.

- In 77% of cases, recommended sunscreens matched those a dermatologist would approve for the participant's skin profile.
- Users appreciated ingredient transparency and ethical flags.

The users expressed their feedback about the user interface design because they wanted to have better control over the weightings which should have included their vegan preference above the product formulation.

The user study showed that participants first identified their skin type incorrectly because they needed the system's visual feedback to fix their wrong beliefs which resulted in increased trust and user participation. The users showed their decision changes when they used the assistant which resulted in 62 percent of users selecting safer sunscreen products while 40 percent of users stopped using their previous comedogenic items, which showed actual changes in their product selection habits.

IX. Qualitative Results and Examples

Figures 3–?? illustrate outputs: facial landmarks and zones, YOLO detections, OCR extraction and a sample recommendation card with rationale.

X. Discussion

The multimodal fusion design improves recommendation relevance and safety by combining visual diagnostics with ingredient-level scrutiny. Important observations:

- Ingredient transparency materially improves user trust.

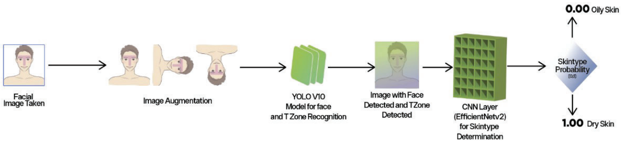

Fig. 4: YOLOv8 acne detection output.

```
=== OCR Extracted Text ===

aquallegica glowt sln screen pa: :na witamin ihen hyj%rete denx

=== Product Found in Database ===

Product Name   : Aqualogica Glow+ Infused Tinted Sunscreen SPF 50+ PA++++ Papaya & Vitamin C(50 g)
Ingredients    : Aloe Vera, Tocopheryl Acetate (Vitamin E), Zinc Oxide, Titanium Dioxide, Glycerin, Butylene Glycol, C12-15 A
lkyl Benzoate, Squalane, Acrylates/C10-30 Alkyl Acrylate Crosspolymer, Centella Asiatica Extract, Panthenol (Vitamin B5), Dis
odium EDTA, Dimethicone, Silica
Suitable Skin  : Combination
Acne Safe      : Yes
```

Fig. 5: OCR extraction.

- Ethical flags are valued and influence decisions for approximately 35% of users in our pilot group.
- On-device inference with TFLite ensures privacy and low latency for core modules; OCR and heavy parsing may run on the server when needed.

The system's explainability functions as an essential element which guides users away from their existing knowledge of brands toward their understanding of actual product components. The proposed method delivers better decision support because it provides clinically accurate information which users can easily understand compared to commercial skincare applications that show unclear results and their components and lack complete ethical disclosure. The solution addresses the typical deficiencies found in standard skincare applications used by consumers.

A. Limitations

- OCR still struggles with reflective foils and heavy embossing.
- Severe dermatological conditions (e.g., cystic acne grade 4) require clinical care and are outside the intended scope.
- CKB coverage requires continuous update as the cosmetics market evolves.

The system provides solid performance yet faces difficulties with fragrance detection because the industry uses inconsistent labeling methods which apply terms such as ”parfum” and ”aroma” and ”fragrance mix.” The identification process for combination skin remains difficult because the oily to normal skin type transition creates unclear boundaries between different skin types. The two problems restrict accurate classification of borderline cases because they show the fundamental uncertainty present in dermatology classification systems.

XI. Ethical Considerations

We emphasise that our assistant is an advisory tool, not a medical device. We implemented:

- explicit consent and optional ephemeral image storage,
- a clear disclaimer recommending dermatological consultation for severe cases,
- provenance metadata for CKB entries and manual curation logs.

The system excludes products which use animal testing and lack clear standards for animal protection. The system only shows brands that possess verifiable cruelty-free certification or public no-testing policies because the system needs to prevent users from accessing products of companies with doubtful ethical standards. The database implements this filtering process which prevents the model from recommending these products through fallback ranking or similarity-based ranking methods.

The system helps users who search for vegan products by marking any product that contains animal-based components and selecting plant-based products as primary options when suitable alternatives exist. The system recognizes all components through ingredient analysis which identifies beeswax lanolin collagen and keratin as non-vegan products unless the manufacturer directly states otherwise. The system maintains recommendation alignment with ethical consumption standards while achieving precise and safe results.

People need to comprehend ingredient analysis because it uses difficult chemical language which leads to incorrect information when statements become too simple. The system uses evidence levels and published dermatology guidelines and known irritation rates instead of making determinative statements to maintain its integrity. The system provides explanations which use neutral language to present accurate information and prevent users from misunderstanding the system as a source of medical guidance.

XII. Conclusion

The proposed system demonstrates how artificial intelligence enables better health decisions through its implementation of dermatological principles into user-friendly health assessment tools. Beyond sunscreen selection, the framework can be extended to monitor skin health trends, detect adverse reactions, and guide users toward safer routines. The increasing public demand for skincare products that provide ethical standards and ingredient information will benefit from AI-driven systems which enable consumers to make better product choices.

XIII. Future Work

Planned extensions:

- Expand to cleansers, moisturisers and serums.
- Build learning-to-rank recommender from user feedback and A/B testing.
- Integrate transformer-based OCR for challenging packages.
- Increase dataset size for clinical-grade acne severity training.

The upcoming developments will implement temporal skin-tracking modules which track skin changes during weeks and months to assess product effects on user skin conditions. The second aspect involves building a small language model which understands dermatology and handles three tasks: it will transform ingredient lists, respond to user questions, and produce basic content. The system will improve user experience through better product matching and customized personal experiences.

Appendix: Dataset Preprocessing and Annotation Protocol

Visual preprocessing: images resized to 512×512 for detection and 224×224 for classification. Photometric augmentation included brightness $\pm 20\%$, contrast $\pm 15\%$ and Gaussian noise $\sigma \in [0, 5]$.

Annotation: Skin type labels were collected via a hybrid approach (self-report + annotator consensus). Lesion boxes required two annotator agreement; disagreements were resolved by a dermatologist. Ingredient canonicalisation used fuzzy matching (Levenshtein threshold 0.85) and manual curation.

Acknowledgements

We thank Asst. Prof. Kundhavai K R and Dr. Sandesh B J for their guidance. This work was conducted as a capstone project at PES University.

References

[1] M. Jayaram et al., "Smart Cosmetics Recommendation System Based on Skin Condition," Industrial Engineering Journal, 2024.
[2] G. M. Rajegowda et al., "AI-Assisted Skincare Routine System in XR," Springer, 2024.
[3] T.-Y. Lin et al., "Facial Skincare Product Recommendation with Computer Vision," Electronics, 2022.
[4] J. Lee et al., "Deep Learning-Based Skincare Product Recommendation," Journal of Cosmetic Dermatology, 2024.
[5] J. Redmon et al., "You Only Look Once: Unified, Real-Time Object Detection," CVPR, 2016.
[6] S. Smith et al., "Robust OCR for Curved Labels Using Synthetic Augmentation," ICDAR, 2022.
[7] Google MediaPipe, "Face Mesh," 2019. [Online]. Available: https://developers.google.com/mediapipe