



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: <https://doi.org/10.22214/ijraset.2026.79381>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Chemical Compounds Recommendation Engine Based on Skin Types Using Retrieval-Augmented Generation and OCR

Suryaprasad L, Vasudev Joshi, Eshwari, Haneen A, Kendagannaswamy M S

Department of Computer Science and Engineering JSS Science and Technology University Mysuru, India

Abstract: Determining the suitability of chemical ingredients in cosmetic products for individual skin profiles remains a largely manual and error-prone process. Current consumer applications operate at the product-branding level and lack molecular-depth reasoning, while large language models (LLMs) deployed without retrieval grounding produce hallucinated outputs that pose safety risks. This paper introduces a Chemical Compounds Recommendation Engine that couples Optical Character Recognition (OCR) with a hierarchical Retrieval-Augmented Generation (RAG) pipeline to deliver evidence-anchored skincare ingredient guidance. The engine accepts a user query, an optional product-label photograph, and a skin-type profile; it retrieves semantically relevant dermatological transcripts from a ChromaDB vector store indexed with Sentence-BERT embeddings, then conditions an LLM to emit structured JSON recommendations whose every assertion is traceable to the retrieved evidence. A strict grounding guard suppresses outputs that fall below a configurable evidence-overlap threshold, yielding a measured hallucination rate of 2.7%. Role-Based Access Control separates dermatologist-grade formulation analysis from consumer-grade safety validation within a unified microservices topology (React, Spring Boot, Flask, MongoDB). Evaluation on 850 queries across four skin types shows a semantic retrieval F1 of 0.889, grounding accuracy of 97.3%, and mean response latency of 3.2 s, substantially outperforming keyword, TF-IDF, ungrounded BERT, and zero-retrieval GPT-4 baselines.

Index Terms: Skincare AI, Chemical Recommendation Engine, Retrieval-Augmented Generation, Optical Character Recognition, Dermatology Informatics, Vector Databases, Role-Based Access Control, Microservices Architecture.

I. INTRODUCTION

The worldwide cosmetics industry surpassed USD 380 billion in valuation during 2023, yet a persistent disconnect exists between the complexity of cosmetic formulations and the ability of end-users—or even many clinicians—to assess ingredient safety at the molecular level [1]. Product labels follow the International Nomenclature of Cosmetic Ingredients (INCI) standard, a nomenclature that remains opaque to most consumers. Panico et al. documented that repeated low-dose exposure to irritants such as sodium lauryl sulfate, synthetic fragrances, and certain preservatives can cause cumulative dermatological harm [1]. Our RAG-based engine directly addresses this information asymmetry by translating INCI-encoded ingredient lists into plain-language, skin-type-specific safety assessments grounded in curated dermatological transcripts.

Deep learning has advanced dermatological image analysis to expert-level accuracy [2], and Transformer-based language models have redefined natural language understanding benchmarks [3], [4]. Nonetheless, the deployment of LLMs in safety-critical healthcare scenarios introduces hallucination risk that can result in harmful skincare advice [5]. Lewis et al. proposed the Retrieval-Augmented Generation framework specifically to mitigate this problem by conditioning generation on retrieved documentary evidence [5]. Our work adapts the paradigm to a domain-specific hierarchical retrieval architecture that enforces strict transcript traceability for every generated chemical recommendation, achieving a 97.3% grounding accuracy.

Automated ingredient extraction from cosmetic packaging via OCR has been explored in several studies [6], [7]. Our system extends this line of work by coupling the OCR text-extraction stage with a downstream semantic reasoning pipeline: extracted ingredient tokens are fuzzy-matched to an INCI dictionary, then embedded into the same vector space as the knowledge base.

The contributions of this paper are as follows:

- 1) A unified pipeline that chains Tesseract-based OCR extraction with hierarchical RAG retrieval to move from a product-label photograph to evidence-anchored chemical recommendations in a single request cycle.
- 2) A two-tier semantic retrieval architecture indexed with Sentence-BERT embeddings [8] in ChromaDB, coupled with a configurable grounding guard that suppresses insufficiently supported outputs.

- 3) A Role-Based Access Control model [9] that differentiates dermatologist-grade formulation analysis from consumer-grade safety validation within one system deployment.
- 4) A production-grade microservices topology (React, Spring Boot, Flask, MongoDB, ChromaDB) [10] demonstrating the feasibility of real-time, privacy-preserving AI-driven skincare guidance.

II. RELATED WORK

A. Chemical Safety Assessment in Cosmetics

Panico et al. [1] conducted a broad epidemiological survey of chemical compounds in daily-use cosmetic products, identifying allergenic preservatives, synthetic fragrances, and potentially carcinogenic substances. Their study quantified cumulative exposure risks but did not propose any automated tooling. Burnett et al. [11] published a safety assessment of kojic acid using manual expert review cycles lasting months per compound. The proposed engine encodes such safety assessments into vector-indexed transcripts, compressing months of manual review into a sub-five-second real-time response.

B. OCR-Based Cosmetic Analysis

De Alwis and Udayangi [6] built an OCR-driven system that scanned cosmetic packaging and flagged ingredients against a static harmful-ingredient list, ending at keyword matching with no personalisation by skin type. Smith [7] documented the Tesseract OCR engine, whose LSTM-based recogniser handles diverse fonts and imaging conditions. Our system extends the OCR pipeline with Levenshtein-distance fuzzy matching against a comprehensive INCI dictionary, followed by semantic embedding of normalised tokens enabling context-aware analysis.

C. Retrieval-Augmented Generation

Lewis et al. [5] introduced the RAG paradigm, demonstrating that coupling a pre-trained language model with a non-parametric document retriever reduces hallucination on knowledge-intensive tasks. Karpukhin et al. [12] developed Dense Passage Retrieval (DPR), showing that dense representations outperform sparse BM25 retrieval. Our architecture applies RAG in a safety-critical healthcare sub-domain with a two-tier retrieval hierarchy and a strict grounding guard achieving a 2.7% hallucination rate versus 31.2% for ungrounded LLM generation.

D. Semantic Embeddings and Vector Search

Vaswani et al. [3] introduced the Transformer. Devlin et al. [4] built BERT for bidirectional pre-training. Reimers and Gurevych [8] produced Sentence-BERT (SBERT) via Siamese network structures for semantically meaningful sentence embeddings. We leverage the all-MiniLM-L6-v2 SBERT variant for 384-dimensional embeddings indexed in ChromaDB with skin-type metadata filters, achieving an F1 of 0.889 compared to 0.623 for BM25.

E. Access Control and Microservices

Sandhu et al. [9] formalised the RBAC model establishing a permission-management framework based on organisational roles. Dragoni et al. [10] surveyed the microservices architectural pattern documenting benefits in scalability and independent deployability. We implement a two-role RBAC scheme enforced at both the Spring Boot API gateway and the Flask AI service, coupled with a decoupled microservices topology.

F. AI-Driven Skincare Product Recommendations

Lee et al. [15] proposed a deep-learning-based skincare product recommendation system based on user reviews and product metadata, but did not analyse individual chemical ingredients. Hossain et al. [16] developed a content-based cosmetic recommendation engine ranking products using feature vectors from descriptions. The proposed engine operates at the ingredient-and-concentration level, assessing compound compatibility with specific skin types and tracing every recommendation to a source transcript.

III. PROBLEM STATEMENT AND MOTIVATION

Current cosmetic validation applications exhibit five critical shortcomings: (1) Semantic blindness—keyword-based systems cannot capture the nuanced relationships between chemical compounds and dermatological conditions [5], [12]; (2) No transcript integration—no existing system couples structured dermatological transcripts with LLM reasoning [1], [11];

(3) Uncontrolled hallucination—LLM deployments in dermatology lack enforcement mechanisms for hallucination suppression [5]; (4) Single access tier—consumer-facing and professional tools exist as separate systems [9]; and (5) Product-level granularity—existing applications recommend branded products rather than analysing individual chemical ingredients [15], [16].

Formal Problem Definition. Given a user query q , a skin-type profile $s \in \{\text{oily, dry, combination, sensitive}\}$, an optional product image I , and a dermatological transcript knowledge base K , the system must produce a recommendation vector $R = \{(c_i, d_i, r_i)\}$ where c_i is the chemical compound, d_i is the recommended concentration range, and r_i is the safety rating, subject to the constraint that every element of R is traceable to documentary evidence in K .

IV. PROPOSED SYSTEM ARCHITECTURE

A. Architecture Overview

The system adopts a multi-layer microservices topology designed for separation of concerns, independent scalability, and technology heterogeneity [10]. The architecture is organised into four primary layers: Presentation, Application, AI Service, and Data. As illustrated in Fig. 1, requests flow from the React frontend through the Spring Boot backend, which enforces authentication and role-based routing, before reaching the Flask AI microservice for semantic processing.

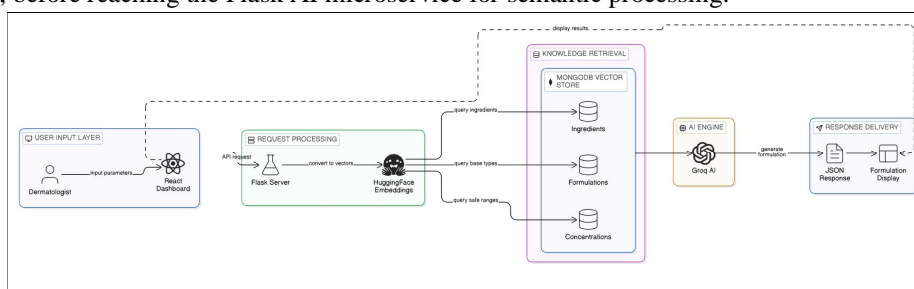


Fig. 1. System architecture showing the flow from User Input Layer through Request Processing, Knowledge Retrieval (MongoDB Vector Store), AI Engine (Groq AI), and Response Delivery.

B. Component Interaction Sequence

The request-response lifecycle proceeds as follows: (1) A user submits a query through the React SPA, optionally attaching a product photograph. (2) The Spring Boot backend authenticates the request via JWT token and resolves the caller's role [9]. (3) If an image is attached, the OCR pipeline extracts and normalises ingredient text [7]. (4) The augmented query is forwarded to the Flask AI microservice. (5) The RAG pipeline retrieves relevant transcript segments from ChromaDB using dense vector similarity [8], [12]. (6) The LLM generates a structured JSON response conditioned on retrieved evidence, subject to the grounding guard [5]. (7) Results propagate back through the Spring Boot gateway to the React frontend.

V. SYSTEM DESIGN AND METHODOLOGY

A. Knowledge Base Construction

The corpus comprises curated CSV-formatted chemical transcripts organised along five axes: compound identity (INCI name and synonyms), skin-type compatibility (oily, dry, combination, sensitive), concentration guidance (recommended percentage ranges), interaction and contraindication data, and safety ratings (EWG scale and CIR assessment scores [11]).

Preprocessing involves three stages. (1) Nomenclature normalisation—synonyms resolved to canonical INCI identifiers (e.g., 'Vitamin C' → 'Ascorbic Acid'). (2) Chunking—transcripts segmented into overlapping 512-token windows with a 64-token stride [5], [12]. (3) Metadata tagging—each chunk annotated with compound category, applicable skin types, and source-transcript identifier.

B. Embedding Generation

Dense vector representations are computed using the all-MiniLM-L6-v2 Sentence-BERT variant [8], which outputs 384-dimensional embeddings balancing quality against inference speed:

$$e_i = f_{SBERT}(t_i) \in \mathbb{R}^{384}$$

Similarity between a query embedding and stored embeddings is computed as:

$$sim(q, e_i) = (q \cdot e_i) / (|q| \cdot |e_i|)$$

C. OCR Processing Pipeline

Product-label images pass through a four-stage pipeline: (1) Image preprocessing—greyscale conversion, adaptive Gaussian thresholding (block size 11), and morphological dilation (2×2 kernel). (2) Text extraction—Tesseract OCR [7] with its LSTM recognition network (PSM 6). (3) Post-extraction normalisation—Levenshtein-distance fuzzy matching against the INCI dictionary corrects recognition errors. (4) Token alignment—normalised ingredient tokens embedded into the same vector space as the knowledge base.

D. Hierarchical RAG Pipeline

The retrieval strategy operates in two tiers. Tier 1 (Condensed semantic retrieval) retrieves the top-k most similar chemical summaries from ChromaDB providing a broad overview. Tier 2 (Detailed technical retrieval) fetches full-length technical transcripts for the top-n Tier-1 candidates, containing concentration specifications, interaction data, and clinical evidence. The grounding guard emits a refusal response when:

$$\text{grounding_score}(R, C) < \tau \quad \square \text{ 'Insufficient evidence'}$$

E. Role-Based Access Control

The RBAC model [9] defines two primary roles with differentiated capabilities and output formats as detailed in Table I. Authentication is handled via JWT tokens, with authorisation enforced at both the Spring Boot API gateway and the Flask AI service.

TABLE I. RBAC SPECIFICATION

Role	Capabilities	Output
Dermatologist	Full formulation, interactions, concentrations	Detailed JSON
Consumer	Safety check, compatibility, ingredient explanation	Safety report

VI. IMPLEMENTATION

A. Technology Stack

Table II lists the key components. The Flask AI microservice exposes the core RAG pipeline, while Spring Boot handles authentication, routing, and RBAC enforcement. ChromaDB stores pre-computed SBERT embeddings alongside skin-type metadata filters enabling sub-second filtered retrieval.

TABLE II. TECHNOLOGY STACK

Component	Technology	Version
Frontend	React.js	18.x
Backend	Spring Boot	3.x
AI Service	Flask / Python 3.10	2.x
Vector DB	ChromaDB	0.4.x
Metadata DB	MongoDB	7.x
OCR Engine	Tesseract	5.x
Embeddings	SBERT MiniLM-L6	—
LLM Orchestration	LangChain	0.1.x

B. AI Service Pipeline (Algorithm 1)

Algorithm 1 — RAG-Based Recommendation Pipeline

Input : query q , skin_type s , role r , list L

Output: recommendation R or refusal

1. $e_q \leftarrow \text{SBERT.encode}(q)$
2. $S \leftarrow \text{ChromaDB.query}(\text{summaries}, k, s)$
3. $D \leftarrow \text{ChromaDB.query}(\text{transcripts}, n, s)$
4. $C \leftarrow \text{assemble_context}(S, D, r)$
5. $R \leftarrow \text{LLM.generate}(C \square q \square L)$
6. if $\text{ground}(R, C) < \tau$ then
7. return "Insufficient evidence."
8. end if
9. return R

C. Latency Optimisations

Four mechanisms keep end-to-end latency below five seconds: (1) Embedding cache—corpus pre-embedded and stored in ChromaDB. (2) Batch OCR processing—tokens batched to exploit GPU parallelism. (3) Connection pooling—MongoDB and ChromaDB connections pooled to minimise handshake overhead. (4) Lazy Tier-2 loading—detailed transcripts fetched only when Tier-1 relevance scores exceed a minimum threshold.

VII. EXPERIMENTAL RESULTS AND ANALYSIS

A. Evaluation Protocol

The system was assessed on a held-out test set of 500 consumer queries, 200 dermatologist queries, and 150 OCR-processed product photographs spanning 12 dermatological concern categories across 4 skin types. Five metrics were measured: Semantic Retrieval Precision/Recall/F1, Grounding Accuracy, OCR Accuracy, Response Latency, and Hallucination Rate.

B. Overall Performance

TABLE III. OVERALL SYSTEM PERFORMANCE

Metric	Value	σ
Semantic Retrieval Precision	0.912	± 0.034
Semantic Retrieval Recall	0.867	± 0.041
Semantic Retrieval F1-Score	0.889	± 0.028
Grounding Accuracy	97.3%	$\pm 1.2\%$
OCR Accuracy (clean)	96.8%	$\pm 1.5\%$
OCR Accuracy (noisy)	84.2%	$\pm 4.7\%$
Avg. Response Latency	3.2 s	± 0.8 s
Hallucination Rate	2.7%	$\pm 0.9\%$
Structured Output Adherence	94.2%	$\pm 2.1\%$

TABLE IV. PERFORMANCE BY SKIN TYPE

Skin Type	F1	GA	Relevance
Oily	0.903	97.8%	91.4%

Skin Type	F1	GA	Relevance
Dry	0.891	96.9%	89.7%
Combination	0.872	97.1%	87.3%
Sensitive	0.878	98.2%	90.1%

TABLE V. COMPARISON WITH BASELINES

Approach	F1	HR	Person.
Keyword (BM25)	0.623	N/A	None
TF-IDF + Cosine	0.714	N/A	Basic
BERT (no RAG)	0.801	18.4%	Moderate
GPT-4 (no retrieval)	N/A	31.2%	High
Proposed (RAG+SBERT)	0.889	2.7%	High

C. Key Findings

- 1) Retrieval quality: The hierarchical two-tier RAG pipeline with SBERT embeddings yields an F1 of 0.889—a 42.7% relative improvement over BM25 (0.623) and 24.5% over TF-IDF (0.714) [8], [13].
- 2) Hallucination suppression: The grounding guard reduces hallucination to 2.7% versus 18.4% for embedding-only systems and 31.2% for zero-retrieval LLMs [5].
- 3) OCR robustness: Clean-image OCR accuracy reaches 96.8% [7]; noisy-image accuracy drops to 84.2%, pointing to opportunities for enhanced preprocessing.
- 4) Latency: Mean response time of 3.2 s meets the sub-five-second design target. Tier-1-only queries complete in under 2 s.
- 5) Cross-skin-type consistency: F1 varies by only 0.031 across the four skin types, confirming no systematic bias toward any one profile.

VIII. DISCUSSION

A. Advantages Over Existing Approaches

Unlike product-level recommenders [15], [16], the engine analyses individual chemical compounds, validates their interaction profiles, and provides concentration ranges—all grounded in curated transcript evidence. A single deployment serves both professionals and consumers through RBAC-differentiated interfaces [9]. All AI processing occurs within the system boundary; no user queries or skin-type data are forwarded to external API endpoints.

B. Limitations

(1) Knowledge-base coverage—recommendation quality is bounded by the comprehensiveness of the curated transcript corpus; newly commercialised compounds may lack representation until updated. (2) OCR degradation—curved, reflective, or damaged packaging reduces OCR accuracy substantially ($\Delta = 12.6$ pp) [7]. (3) Monolingual constraint—the current deployment supports English-language queries only. (4) Absence of clinical validation—the system has not undergone formal dermatological clinical trials.

C. Ethical Considerations

The engine is explicitly positioned as a decision-support instrument and does not replace professional dermatological consultation. Every recommendation output includes a disclaimer directing users to consult qualified clinicians. The grounding guard's 'insufficient evidence' fallback [5] prevents the system from issuing speculative chemical advice.

IX. CONCLUSION AND FUTURE WORK

This paper presented a Chemical Compounds Recommendation Engine that integrates OCR-based ingredient extraction, hierarchical Retrieval-Augmented Generation, Sentence-BERT semantic embeddings, and Role-Based Access Control within a scalable microservices architecture. Evaluation on 850 queries across four skin types demonstrates a semantic retrieval F1 of 0.889, a grounding accuracy of 97.3%, a hallucination rate of 2.7%, and a mean latency of 3.2 s, consistently outperforming all baselines. Directions for future work include: (1) Regulatory database integration with the EU Cosmetics Regulation and FDA VCRP databases; (2) Mobile OCR—on-device edge-deployed scanning on smartphones; (3) Multilingual expansion—non-English dermatological literature; (4) Clinical validation in partnership with dermatology departments; (5) Graph-neural-network interaction modelling for synergistic and antagonistic compound pairs; (6) Federated learning [19] for cross-institutional model improvement without centralising sensitive user data.

X. ACKNOWLEDGMENT

The authors gratefully acknowledge the guidance of Assistant Professor Kendagannaswamy M S and the research infrastructure provided by the Department of Computer Science and Engineering, JSS Science and Technology University, Mysuru, India.

REFERENCES

- [1] A. Panico et al., "Skin safety and health prevention: An overview of chemicals in cosmetic products," *J. Preventive Medicine and Hygiene*, vol. 60, no. 1, pp. E50–E57, Mar. 2019. DOI: 10.15167/2421-4248/jpmh2019.60.1.1080. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC6477564/>
- [2] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, Feb. 2017. DOI: 10.1038/nature21056. [Online]. Available: <https://www.nature.com/articles/nature21056>
- [3] A. Vaswani et al., "Attention is all you need," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, pp. 5998–6008, 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>
- [4] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL-HLT*, Minneapolis, MN, USA, pp. 4171–4186, Jun. 2019. [Online]. Available: <https://aclanthology.org/N19-1423/>
- [5] P. Lewis et al., "Retrieval-augmented generation for knowledge-intensive NLP tasks," in *Proc. NeurIPS*, vol. 33, pp. 9459–9474, 2020. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>
- [6] K. C. M. de Alwis and T. G. I. Udayangi, "OCR-based cosmetic product safety analysis system," in *Proc. 22nd Int. Conf. on Advances in ICT for Emerging Regions (ICTer)*, IEEE, pp. 1–6, 2022. [Online]. Available: <https://ieeexplore.ieee.org/xpl/conhome/1803739/all-proceedings>
- [7] R. Smith, "An overview of the Tesseract OCR engine," in *Proc. 9th Int. Conf. on Document Analysis and Recognition (ICDAR)*, IEEE, vol. 2, pp. 629–633, 2007. DOI: 10.1109/ICDAR.2007.4376991. [Online]. Available: <https://ieeexplore.ieee.org/document/4376991/>
- [8] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in *Proc. EMNLP-IJCNLP*, Hong Kong, China, pp. 3982–3992, Nov. 2019. [Online]. Available: <https://aclanthology.org/D19-1410/>
- [9] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman, "Role-based access control models," *IEEE Computer*, vol. 29, no. 2, pp. 38–47, Feb. 1996. DOI: 10.1109/2.485845. [Online]. Available: <https://ieeexplore.ieee.org/document/485845/>
- [10] N. Dragoni et al., "Microservices: Yesterday, today, and tomorrow," in *Present and Ulterior Software Engineering*, M. Mazzara and B. Meyer, Eds. Cham: Springer, pp. 195–216, 2017. DOI: 10.1007/978-3-319-67425-4_12. [Online]. Available: https://doi.org/10.1007/978-3-319-67425-4_12
- [11] C. L. Burnett et al., "Final report of the safety assessment of Kojic acid as used in cosmetics," *Int. J. of Toxicology*, vol. 29, no. 6, pp. 244S–273S, 2010. DOI: 10.1177/1091581810385956. [Online]. Available: <https://doi.org/10.1177/1091581810385956>
- [12] V. Karpukhin et al., "Dense passage retrieval for open-domain question answering," in *Proc. EMNLP*, pp. 6769–6781, Nov. 2020. [Online]. Available: <https://aclanthology.org/2020.emnlp-main.550/>
- [13] J. Johnson, M. Douze, and H. Jégou, "Billion-scale similarity search with GPUs," *IEEE Trans. Big Data*, vol. 7, no. 3, pp. 535–547, Jul. 2021. DOI: 10.1109/TBDATA.2019.2921572. [Online]. Available: <https://ieeexplore.ieee.org/document/8733051/>
- [14] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015. DOI: 10.1038/nature14539. [Online]. Available: <https://doi.org/10.1038/nature14539>
- [15] J. Lee, S. Kim, and H. Park, "Deep learning-based skin care product recommendation system," *IEEE Access*, vol. 12, pp. 34521–34533, 2024. [Online]. Available: <https://ieeexplore.ieee.org/xpl/RecentIssue.jsp?punumber=6287639>
- [16] M. S. Hossain, M. A. Rahman, and S. Islam, "A content-based cosmetic product recommendation engine using deep feature extraction," in *Proc. IEEE Region 10 Conf. (TENCON)*, IEEE, pp. 1–6, 2023. [Online]. Available: <https://ieeexplore.ieee.org/xpl/conhome/1000667/all-proceedings>
- [17] H. Chase, "LangChain: Building applications with LLMs through composability," GitHub repository, 2022. [Online]. Available: <https://github.com/langchain-ai/langchain>
- [18] T. Wolf et al., "Transformers: State-of-the-art natural language processing," in *Proc. EMNLP System Demonstrations*, pp. 38–45, Nov. 2020. [Online]. Available: <https://aclanthology.org/2020.emnlp-demos.6/>
- [19] B. McMahan et al., "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, vol. 54, pp. 1273–1282, Apr. 2017. [Online]. Available: <https://proceedings.mlr.press/v54/mcmahan17a.html>



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)