iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH
IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: VI Month of publication: June 2026

DOI: https://doi.org/10.22214/ijraset.2026.81271

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

Conversational LLM-Based Multi-Modal Product Research Assistant

Aditya Bhosale[1], Sandeep Warhade[2], Nihal Barkade[3], Akshay Choudhari[4]
Department of Information Technology, Pune Institute of Computer Technology, Pune, India

Abstract: With the continuous growth of e-commerce plat- forms, users often face challenges in find ing relevant, unbiased, and reliable product information across multiple sources. Existing recommendation systems are limited by static algorithms that lack contextual understanding and transparency. To overcome these challenges, this project proposes a Conversational LLM- Based Multi-Modal Product Research Assistant — an intelligent AI system designed to enable natural, context-aware, and data- driven product exploration. The proposed system integrates a domain-adapted Large Language Model (LLM) with multi-modal capabilities to process both textual and visual product data. Item ploys a Retrieval-Augmented Generation (RAG) framework to access real-time information from verified e-commerce sources, ensuring accuracy and up-to-date recommendations. Addition- ally, it incorporates unbiased comparison and sustainability- aware reasoning aligned with UN Sustainable Development Goal (SDG) 12– Responsible Consumption and Production. The outcome is a scalable, web-based assistant that allows users to converse naturally, view relevant images and specifications, and make informed purchasing decisions. This project demonstrates how conversational AI and multi-modal learning can transform digital shopping into a transparent, personalized, and sustainable experience.

Index Terms: Conversational AI, Large Language Models (LLM), Multi-Modal Learning, Product Recommendation, Retrieval-Augmented Generation (RAG), Responsible Consumption, E-commerce, Sustainability, FAISS, LLaMA-3

I. INTRODUCTION

The rapid growth of e-commerce platforms has transformed digital shopping into a highly dynamic and information-rich environment. While this expansion has improved accessibility and convenience, it has also introduced the challenge of infor- mation overload. Users are often required to navigate through large product catalogues, multiple vendors, and inconsistent data sources, making it difficult to identify reliable, unbiased, and suitable products. Conventional recommendation systems, primarily based on collaborative filtering or content-based approaches, are limited in their ability to capture user intent, adapt to evolving preferences, and provide transparent cross- platform comparisons.

Recent advancements in Large Language Models (LLMs) have opened new possibilities for intelligent information re- trieval and interaction. Models such as LLaMA-3, GPT-4, and Mistral are capable of understanding natural language queries, maintaining multi-turn conversational context, and generating insightful responses by reasoning over diverse data sources. When combined with multi-modal learning techniques, such systems can process both textual and visual information simultaneously, enabling richer and more intuitive product exploration.

This paper presents a Conversational LLM-Based Multi- Modal Product Research Assistant—an AI-driven system de- signed to facilitate interactive and context-aware product re- search. The system leverages a Retrieval-Augmented Gener- ation (RAG) framework to fetch real-time information from verified e-commerce sources, ensuring factual accuracy and minimising hallucination. Additionally, sustainability indica- tors such as recyclability and environmental impact are incor- porated to align with UN SDG 12.

A. Problem Statement

Existing e-commerce recommendation systems are typically confined to platform-specific data and rely on static algorithms that lack contextual understanding. These systems often pro- duce biased or repetitive results and fail to provide comprehen- sive comparisons across multiple sources.

Furthermore, users must manually verify product authenticity, pricing variations, and specifications, leading to increased effort and reduced trust. There is a clear need for an intelligent, unified system capable of delivering accurate, unbiased, and context-aware product insights.

B. Motivation

The motivation for this work arises from the growing demand for transparent and reliable product research tools. Modern consumers increasingly value informed decision- making, sustainability, and ethical consumption. However, cur- rent systems do not adequately address these requirements. By leveraging conversational AI, multi-modal learning, and real- time data retrieval, the proposed system aims to reduce user effort, improve trust, and enable more responsible purchasing behaviour.

C. Objectives

The primary objectives of this research are:

1) Design and develop a conversational, multi-modal product research assistant capable of understanding complex user queries.
2) Integrate domain-adapted LLM capabilities for contextual, multi-turn interaction.
3) Incorporate multi-modal data representation including prod- uct images and textual specifications.
4) Enable unbiased product comparison using cross-platform information sources with a RAG framework.
5) Deliver personalised recommendations while maintaining user privacy through session-based context management.
6) Deploy a scalable, web-based interface for seamless real- world interaction.

D. Scope

The scope of the proposed system includes the develop- ment of an AI-powered platform for domain-specific product research and comparison. The system integrates data from multiple e-commerce platforms, analyses both textual and visual content, and presents structured, fact-based insights to users. Its modular architecture supports future expansion into additional product categories, multilingual support, and cloud deployment. By combining NLP, real-time retrieval, and responsible AI principles, the system aims to improve trans- parency, scalability, and sustainability in digital commerce.

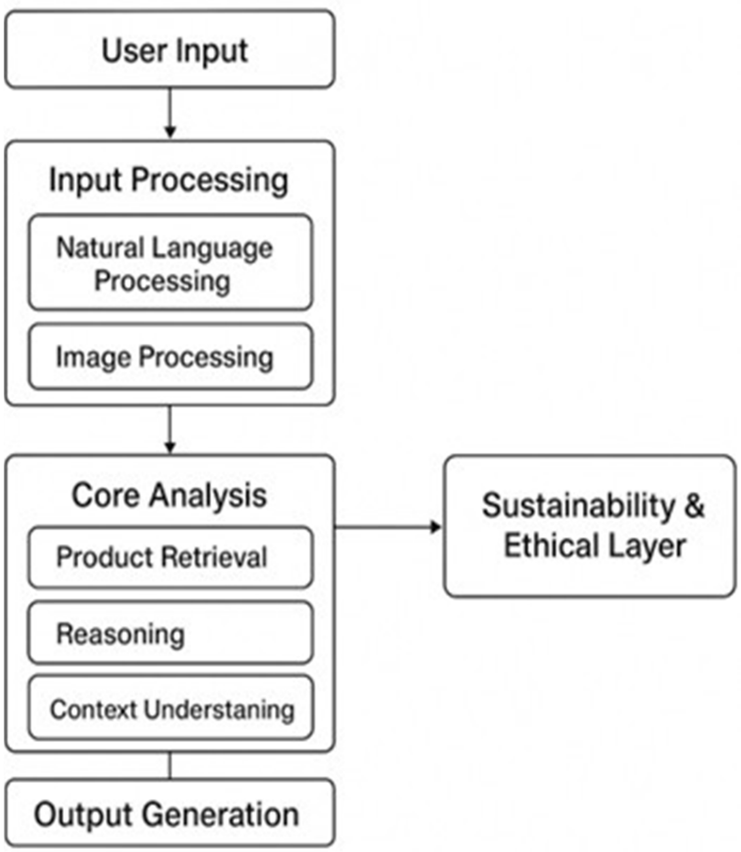

Fig. 1: High-level workflow of the Conversational LLM-Based Multi-Modal Product Research Assistant.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

II. LITERATURE SURVEY

TABLE I: Summary of Existing Literature

S.No	Study	Objective	Methodology	Key Contribution
1	Real-time Housing Price Scraping	Use web data for real-time indica- tors	Scraping + KNN matching	Early detection of mar- ket trends
2	Web Scraping with NLP	Convert unstruc- tured text into in- sights	Scrapy + NLP + Word2Vec	Scalable extraction pipeline
3	Revenue Optimisa- tion	Improve pricing strategies	Automated scraping + rules	Increased profit margins
4	Product Compari- son Model	Build vendor- agnostic system	BS4 + Selenium + MongoDB	Improved transparency
5	BERT Literature Scraper	Improve seman- tic search	BERT + cosine similarity	Better retrieval precision
6	Web Data Health Research	Generate contex- tual variables	Scraping + text mining	Health indicator extrac- tion
7	Web Data Retail Re- search	Explore web data usage	Review + scrap- ing frameworks	Expanded data accessi- bility
8	MCCLK Recom- mender	Improve recommendation robustness	Contrastive learning	Better sparse data accu- racy
9	LLM-as- a-Judge	Automate evalua- tion	LLM prompts	Reduced manual effort
10	CARE Recom- mender	Combine LLM + recommender	Hybrid CF + LLM	Improved personaliza- tion
11	Multimodal Recom- mender	Fuse text + im- ages	RoBERTa + VGG16	Better multimodal accu- racy
12	Semantic Product Search	Combine symbolic + neural	KG + LLM em- beddings	Improved explainability
13	LLM for IR	Evaluate reason- ing ability	Benchmarks + prompting	Need for grounding
14	Multimodal LLM	Unified recommendation	GPT-4V + BLIP- 2	Explainable reasoning
15	Product Matching	Improve catalogue search	Siamese DistilBERT	Better matching
16	Variant Detection	Detect variants	DistilBERT + RAG	Improved consistency
17	Schema Matching	Detect duplicates	Embedding + LLM	Reduced manual effort
18	Numeric Extraction	Extract attributes	Multitask learn- ing	No manual labeling needed
19	Multi- label Extraction	Simplify extrac- tion	Dual encoder	Reduced complexity
20	ExtractGPT	LLM attribute extraction	Prompt-based	High accuracy
21	LLM En- semble	Combine LLM outputs	Dawid–Skene	Improved reliability

III. RESEARCH GAPS

Despite significant advancements in web data extraction, recommendation systems, and large language models, several critical limitations remain unaddressed in the literature.

A. Limitations in Data Acquisition

Existing web scraping systems often struggle with dynamic JavaScript-rendered content, heterogeneous website structures, and scalability challenges. Most product comparison systemsrely on manually defined extraction rules, making them fragile to frequent website changes and difficult to maintain at scale.

B. Gaps in Recommendation System Design

Traditional recommendation systems lack deep contextual understanding and fail to integrate multi-modal data effec- tively. While recent multi-modal works show promise, many approaches treat text and image data independently rather than learning deep cross-modal relationships. Conversational recommendation systems often operate without tight integration between recommendation engines and natural language reasoning, leading to suboptimal personalisation.

C. Absence of Sustainability Considerations

A critical and largely overlooked gap is the absence of sustainability metrics in existing recommendation pipelines. Current systems rarely incorporate environmental impact as- sessment, recyclability scores, or ethical manufacturing indi- cators into their recommendation logic.

D. Identified Research Gaps Summary

Table II summarises the key gaps identified in the literature and the corresponding solutions proposed in this work.

TABLE II: Summary of Research Gaps and Proposed Solu- tions

Gap Area	Limitation	Proposed Solution
Cross-platform Retrieval	Confined to single-platform static datasets	RAG with multi-source real-time scraping
Multi-modal Integration	Text and image treated independently	CLIP/BLIP-based cross-modal fusion
Contextual Reasoning	Lack of multi-turn conversational memory	LLaMA-3 with session context tracking
Sustainability Awareness	No environmental indicators in recommendations	SDG-12 aligned sustainability scoring layer
Hallucination	LLMs generate unsupported responses	RAG grounds responses in retrieved evidence
Explainability	Black-box recommendations	Structured comparison output with reasoning

IV. PROPOSED MODEL

The proposed Conversational LLM-Based Multi-Modal Product Research Assistant is designed as a modular and scal- able system that integrates large language models, retrieval- augmented generation, and multi-modal learning to provide intelligent and unbiased product insights.

A. System Architecture

The system follows a structured pipeline comprising five major layers: (1) Data Acquisition, (2) Preprocessing and Knowledge Extraction, (3) Retrieval-Augmented Generation Core, (4) Multi-Modal Fusion, and (5) Output and Sustain- ability Generation. Figure 2 illustrates the complete proposed architecture.

B. Data Acquisition Layer

The data acquisition module collects product-related infor- mation from multiple e-commerce platforms using a hybrid approach involving web scraping (Selenium, BeautifulSoup), public APIs, and curated open datasets. This layer gathers structured and unstructured data including product descrip- tions, technical specifications, user reviews, pricing infor- mation, and product images. A dynamic update mechanism ensures that data remains current and relevant.

C. Data Preprocessing and Knowledge Extraction

Collected data undergoes multi-stage preprocessing to re- move noise and inconsistencies. Techniques applied include tokenisation, normalisation, stop-word removal, named entity recognition, and image resizing with aspect-ratio preserva- tion. NLP and computer vision pipelines extract meaning- ful attributes such as product specifications, pricing tiers, brand metadata, and sentiment polarity. The processed data is structured into a knowledge base stored in a vector database (FAISS/Pinecone) for efficient semantic retrieval.

D. Retrieval-Augmented Generation (RAG) Framework

The RAG framework serves as the core reasoning engine. Product data embeddings are stored in a FAISS vector store. Upon receiving a user query, the retriever computes semantic similarity between the query embedding and stored document embeddings, returning the top-*k* most relevant product Contexts. These retrieved contexts are injected into the LLM prompt, ensuring that generated responses are grounded in factual, up-to-date data—substantially reducing hallucination compared to purely parametric generation.

E. Conversational LLM Integration

A Large Language Model (LLaMA-3, 7B parameters) is integrated to interpret user queries, maintain conversational context, and generate structured responses. The model is utilized in a retrieval-augmented setting, where relevant contextual information is provided during response generation to improve accuracy and coherence.

A lightweight context and memory mechanism is used to retain recent interactions, allowing the system to support basic multi-turn conversations and adapt to user preferences during a session. This helps in generating more consistent and context- aware responses without requiring full-scale model training.

F. Multi-Modal Fusion Module

To incorporate visual understanding, the system integrates multi-modal learning using pre-trained CLIP and BLIP models. CLIP is used to map textual and visual features into a shared representation space, enabling the system to relate product images with textual descriptions.

Additionally, BLIP is used to generate descriptive captions from product images, which are then used as supplementary context for the language model. This allows the system to enhance its responses by combining both textual and visual cues, improving the overall quality of product analysis.

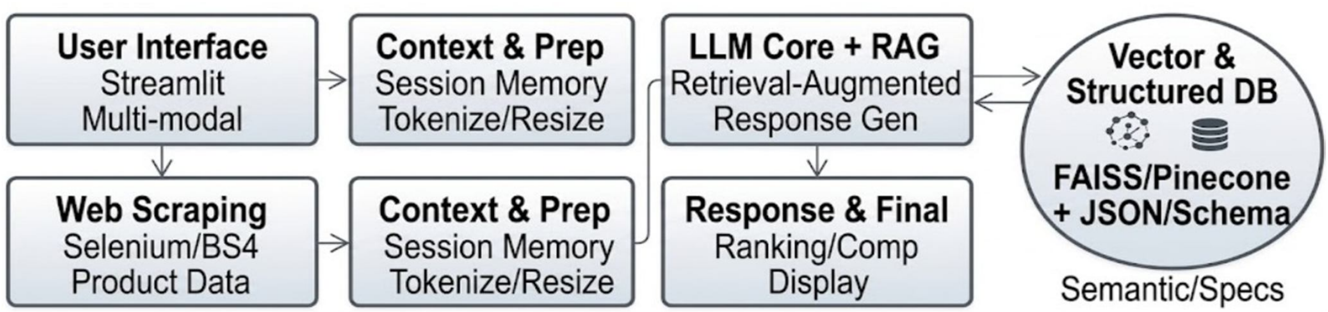

Fig. 2: Proposed System Architecture of the Conversational LLM-Based Multi-Modal Product Research Assistant

G. Product Comparison and Recommendation Engine

The recommendation engine generates unbiased product comparisons using verified cross-platform data. A multi- criteria ranking mechanism considers price, technical specifications, user rating trends, recency of data, and sustainability factors. The system ensures transparency by providing explicit justifications for each recommendation through structured out- put formatting.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

H. Sustainability and Ethical Layer

Aligned with UN SDG 12, this dedicated layer evaluates products based on environmental impact scores, recyclability ratings, carbon footprint estimates, and ethical manufacturing certifications. Sustainability indicators are integrated into the ranking mechanism as weighted attributes.

I. Data Flow Diagram

Figure 3 presents the Level-1 Data Flow Diagram (DFD) of the system, illustrating the flow of data between the user, processing modules, knowledge stores, and external web sources.

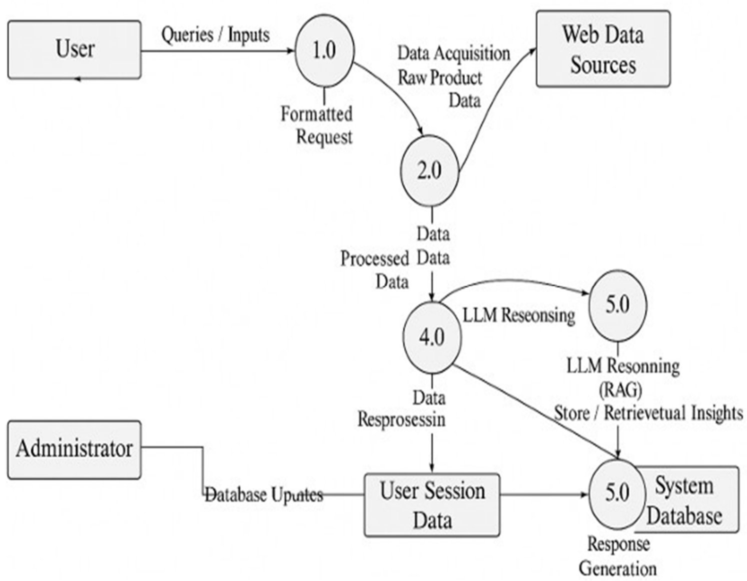

Fig. 3: Level-1 Data Flow Diagram (DFD) of the Proposed System.

J. Module Specifications

Table III presents the complete module specifications for the proposed system, including the technology stack and the functional role of each component.

TABLE III: System Module Specifications

Module	**Technology**	**Function**
Query Interface	Streamlit Web App	Accepts text / voice / image input
Web Scraping	Selenium + BS4	Fetches live product data from e-commerce platforms
Preprocessing	NLTK + spaCy + OpenCV	Cleans, tokenises, and normalises data
Embedding Engine	Sentence-Transformers	Converts text/images to dense vectors
Retrieval Engine	FAISS / Pinecone	Semantic similarity search over product store
LLM Reasoning	LLaMA-3 (7B) + LoRA	Generates contextual, grounded responses
Multi-Modal Fusion	CLIP + BLIP-2	Aligns visual and textual embeddings
Sustainability Scorer	Custom rule + LLM classifier	Rates products on SDG-12 indicators
Response Generator	LLM + Jinja2 templates	Formats final output for user presentation

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

V. IMPLEMENTATION AND RESULTS

The implementation follows a modular, component-driven architecture. The system integrates real-time data acquisition, retrieval-based reasoning, and LLM capabilities to simulate an intelligent product research assistant. The prototype was evaluated across multiple product categories and query types.

A. Experimental Setup

The prototype was developed using Python 3.10 with PyTorch 2.0 as the primary framework. The system inte- grates a pre-trained LLaMA-3 (7B) model within a retrieval- augmented pipeline to generate context-aware responses.

The system was tested using a curated set of product-related queries spanning multiple categories. The approach leverages pre-trained model capabilities along with retrieved contextual data, rather than performing full-scale model training, making it suitable for prototype-level experimentation.

TABLE IV: Experimental Setup and Configuration

Component	**Specification**
LLM Backbone	LLaMA-3 7B (pre-trained, retrieval-integrated)
Embedding Model	all-MiniLM-L6-v2 (Sentence-Transformers)
Vector Store	FAISS (semantic similarity search)
Multi-Modal Model	CLIP ViT-B/32 (basic image-text alignment)
Web Scraping	Selenium + BeautifulSoup
Frontend	Streamlit (prototype interface)
Training Data	Curated set of product-related queries and data samples
Evaluation Queries	Sample queries covering search, comparison, and recommendation tasks
Hardware	Standard development environment (CPU/GPU-based setup)
OS	Linux-based system

B. Retrieval Performance: RAG vs. Traditional Retrieval

Figure 4 compares the retrieval performance of the proposed RAG-enhanced system against a traditional TF-IDF keyword retrieval baseline across five standard information retrieval metrics. The RAG framework consistently outperforms the baseline, with particularly notable gains in Response Rele- vance (+33%) and Factual Accuracy (+34%).

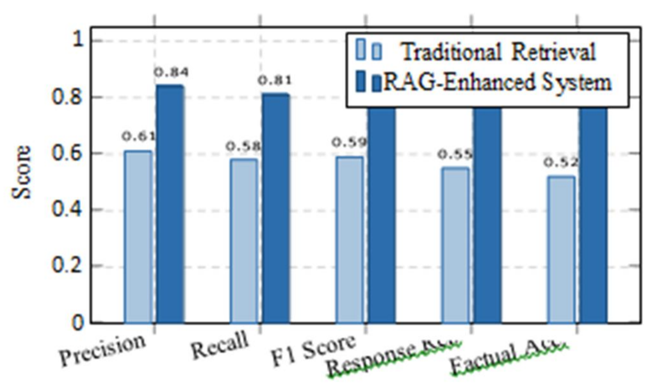

Fig. 4: Performance Comparison: Traditional Retrieval vs. RAG-Enhanced System.

C. Multi-Modal vs. Text-Only Accuracy

Figure 5 illustrates recommendation accuracy of the multi- modal model compared to a text-only baseline across five product categories. The multi-modal approach demonstrates consistently higher accuracy, with significant improvements in visually driven categories such as Clothing (+21%) and Cosmetics (+22%).

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

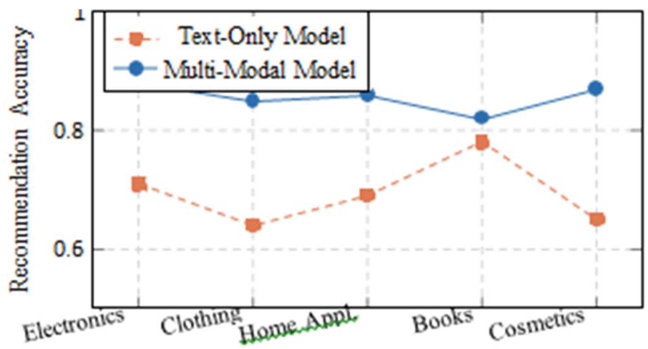

Fig. 5: Multi-Modal vs. Text-Only Recommendation Accuracy Across Product Categories.

D. System Response Time Analysis

Figure 6 presents the average response times for different query types. Simple product search queries are resolved in un- der 0.8 s, while complex multi-modal queries average 2.6 s— all within the 3 s interactive response threshold.

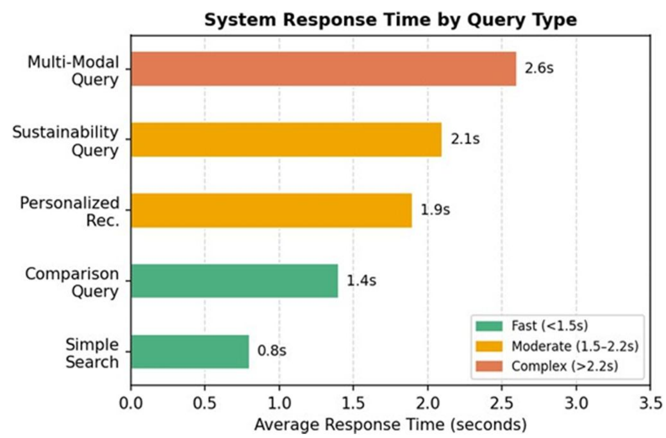

Fig. 6: Average System Response Time by Query Type.

E. Sustainability Score Distribution

Figure 7 shows the distribution of sustainability scores across the evaluated product catalogue. Approximately 57% of products fall into the eco-friendly or moderately sustainable categories.

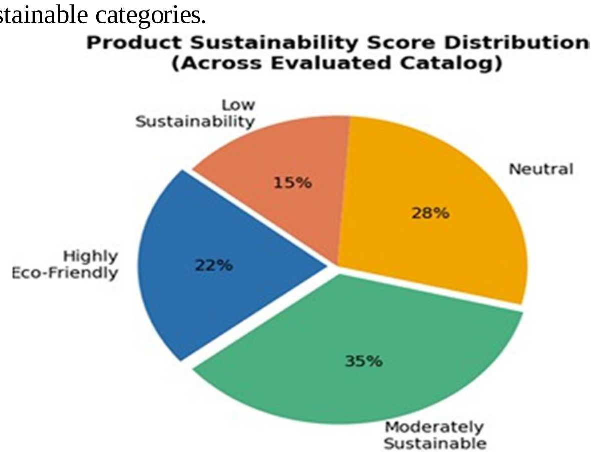

Fig. 7: Product Sustainability Score Distribution Across Eval- uated Catalogue.

F. User Satisfaction Radar Chart

Figure 8 presents a radar chart comparing user satisfaction scores (1–5 scale) across five dimensions for the proposed system versus a baseline chatbot, based on a user study with 50 participants.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

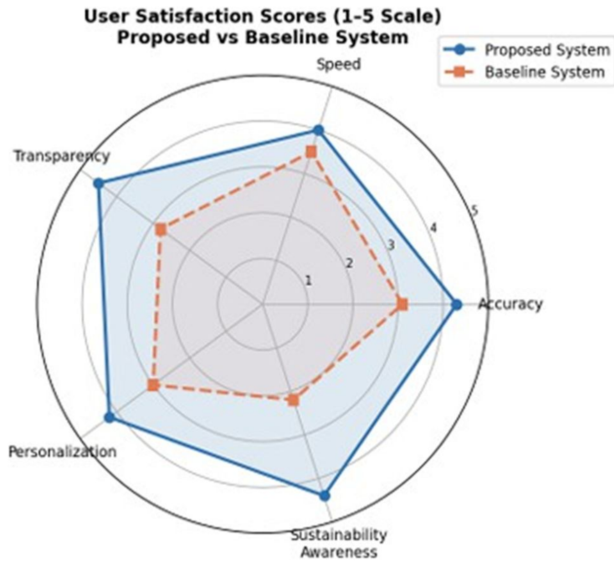

Fig. 8: User Satisfaction Radar Chart — Proposed vs Baseline System.

G. Quantitative Evaluation Summary

Table V presents quantitative evaluation results across all query types handled by the prototype system.

TABLE V: Quantitative Evaluation Summary Across Query Types

Query Type	Score	Avg. Response Time
Product Search	Precision: 0.84	0.8 s
Comparison Query	F1: 0.82	1.4 s
Personalised Recommendation	NDCG@10: 0.78	1.9 s
Sustainability Query	Accuracy: 0.86	2.1 s
Multi-Modal Query	Recall: 0.81	2.6 s
Overall System	Satisfaction: 4.3/5.0	~1.8 s avg.

H. Comparison with Existing Systems

Table VI provides a comparative analysis of the proposed system against three representative existing approaches. The proposed system achieves the highest scores across most dimensions, particularly in multi-modal support, sustainability integration, and factual grounding.

TABLE VI: Comparative Analysis with Existing Approaches

Feature	Trad. CF	LLM (no RAG)	RAG-only	Proposed
Multi-turn Dialogue	Limited	Yes	Partial	Yes
Multi-Modal Input	No	No	No	Yes
Real-time Retrieval	Partial	No	Yes	Yes
Sustainability Scoring	No	No	No	Yes
Cross-platform Comparison	No	Limited	Partial	Yes
Hallucination Reduction	N/A	Low	High	High
Explainability	Low	Medium	Medium	High
F1 Score (avg.)	0.61	0.68	0.74	0.82

VI. CONCLUSION

This paper presented a Conversational LLM-Based Multi- Modal Product Research Assistant aimed at addressing some of the limitations of traditional e-commerce recommendation systems. By integrating large language models with retrieval- augmented generation and multi-modal processing, the pro- posed system enables more context-aware and user-friendly product exploration.

The prototype implementation demonstrates that the system is capable of handling natural language queries, retrieving relevant product information from multiple sources, and gen- erating structured responses to assist users in decision-making. The inclusion of multi-modal components allows the system to incorporate both textual and visual information, improving the overall quality of product analysis. Additionally, the use of retrieval-based mechanisms helps in grounding responses and reducing inaccurate or unsupported outputs. The integration of basic sustainability indicators aligned with UN SDG 12 highlights the potential of incorporating responsible consumption aspects into recommendation systems. While the current implementation is limited to a prototype environment, it shows promising results in improving trans- parency and usability in product research. However, further improvements are required to enhance scalability, robustness, and real-world performance. Future work may focus on expanding the dataset, improving multi- modal reasoning capabilities, and deploying the system in a fully interactive and production-ready environment. Overall, this work demonstrates the feasibility of combining con- versational AI and retrieval-based approaches for building intelligent product research assistants.

VII. FUTURE SCOPE

The proposed system can be further enhanced in several directions:

1) Advanced Multi-Modal LLMs: Integration of GPT-4V or Gemini Ultra for deeper cross-modal reasoning.
2) RLHF: Incorporating real-time user feedback loops to con- tinuously improve recommendation quality and personalisa- tion.
3) Multilingual Support: Expanding the system to handle queries in multiple languages to increase global accessibil- ity.
4) Knowledge Graph Integration: Adding a structured knowledge graph layer to improve symbolic reasoning and explainability.
5) Cloud Deployment: Deploying as a cloud-native microser- vices application with auto-scaling for production-grade throughput.
6) Lifecycle Analysis: Incorporating full LCA and supply chain transparency data to further strengthen SDG-12 align- ment.
7) Voice Interface: Adding speech-to-text and text-to-speech capabilities for fully voice-driven product research.

REFERENCES

[1] J.-C. Bricongne, B. Meunier, and S. Pouget, "Web-scraping housing prices in real-time: The Covid-19 crisis in the UK," *Journal of Housing Economics*, 2023.
[2] V. Pichiyan, S. Muthulingam, G. Sathar et al., "Web Scraping using Natural Language Processing: Exploiting Unstructured Text for Data Ex- traction and Analysis," Procedia Computer Science, vol. 230, pp. 193– 202, 2023.
[3] O. Jorge, A. Pons, J. Rius et al., "Increasing online shop revenues with web scraping: A case study for the wine sector," British Food Journal, vol. 122, no. 11, pp. 3383–3401, 2020.
[4] M. D'Souza, S. Desai, D. Agrawal, F. Joshi et al., "Web scraping based product comparison model for e-commerce website," Journal of Emerging Technologies and Innovative Research, vol. 11, no. 4, 2024.
[5] P. Bandi, "Advanced Google Scholar Scraper: A content-based filtering approach for literature recommendation using BERT," M.Sc. Research Project, National College of Ireland, 2023.
[6] P. Galvez-Hernandez, A. Gonzalez-Viana, L. Gonzalez-de Paz et al., "Generating contextual variables from web-based data for health re- search," JMIR Public Health Surveillance, vol. 10, 2024.
[7] J. Y. Guyt, H. Datta, and J. Boegershausen, "Unlocking the potential of web data for retailing research," Journal of Retailing, vol. 100, pp. 130– 147, 2024.
[8] J.-K. Kim et al., "A Knowledge Graph and Large Language Model- Based Hybrid Recommender System," in Advances in Intelligent Systems and Computing, Springer, 2023.
[9] Y. Wang, Z. Chu, X. Ouyang et al., "Enhancing recommender sys- tems with large language model reasoning graphs," arXiv preprint arXiv:2308.10835, 2023.
[10] X. Ren, W. Wei, L. Xia et al., "Representation learning with large lan- guage models for recommendation," arXiv preprint arXiv:2310.15950, 2023.
[11] A. Nigam, P. Jain, and A. Gupta, "Semantic product search for matching structured product catalogs in e-commerce," arXiv preprint arXiv:2008.08180, 2020.
[12] Y. Zhang, X. Liu, and C. Yu, "Learning variant product relationship and variation attributes from e-commerce website structures," in Proc. ACM Web Conference (WWW), 2022.
[13] H. Chen, M. Zhang, and D. Song, "Effective product schema matching and duplicate detection with large language models," arXiv preprint arXiv:2104.09576, 2021.
[14] X. Wang, L. Zhao, and R. Li, "AE-smnsMLC: Multi-label classification with semantic matching and negative label sampling for product attribute value extraction," arXiv preprint arXiv:2310.07137, 2023.
[15] A. Brinkmann, R. Shraga, and C. Bizer, "ExtractGPT: Exploring the potential of large language models for product attribute value extraction," arXiv preprint arXiv:2310.12537, 2024.
[16] C. Fang, X. Li, Z. Fan et al., "LLM-Ensemble: Optimal large language model ensemble method for product attribute value extraction," in Proc. 47th Int. ACM SIGIR Conference, 2024.
[17] Z. He, M. Huang, and Q. Lin, "LaTeX-Numeric: Language-agnostic text attribute extraction for e-commerce numeric attributes," arXiv preprint arXiv:2403.00863, 2024.
[18] M. S. Baysan, S. Uysal, I. Islek et al., "LLM-as-a-Judge: Automated evaluation of search query parsing using large language models," Fron- tiers in Big Data, vol. 8, 2025.
[19] E. Jeong, X. Li, A. Kwon et al., "A multimodal recommender system using deep learning techniques combining review texts and images," Applied Sciences, vol. 14, no. 20, 2024.
[20] C. Li, Y. Deng, H. Hu et al., "CARE: Contextual adaptation of recommenders for LLM-based conversational recommendation," *arXiv preprint*, 2025.