Semantic Search and Multimodal Personalization in Fashion E-Commerce Using Large Language Models and Vector Databases

[1]Priyadharshini M, [2]Gokul Nath R, [3]Shahul Hameed S, [4]Pavithran A

1 Assistant Professor, Department of Computer Science and Engineering, SRM Valliammai Engineering College, Tamil Nadu, India

2,3,4 Final Year Student, Department of Computer Science and Engineering, SRM Valliammai Engineering College, Tamil Nadu, India

Abstract : Online fashion shopping often fails to deliver accurate results due to the limitations of traditional keyword-based search systems. Users typically describe their needs using natural language, cultural terms, or occasion-based queries, which are not effectively handled by existing platforms. This paper presents FashionAI, a full-stack AI-powered fashion recommendation system that understands user queries in plain language and generates personalized outfit suggestions. The system integrates a large language model for query parsing and multimodal analysis, semantic vector search using sentence embeddings, and a rule-based engine for outfit generation and colour compatibility. It supports both text-based and image-based inputs, enabling the extraction of user attributes such as gender, age group, skin tone, and body type for enhanced personalization. A multi-signal ranking mechanism combines semantic similarity, occasion relevance, user preferences, and visual attributes to improve recommendation quality. The system also incorporates cultural and seasonal awareness to better match real-world fashion needs. Experimental evaluation across multiple scenarios demonstrates high accuracy in occasion detection, correct gender mapping, and consistent outfit generation, with response times under a few seconds. The results indicate that integrating language models, semantic search, and rule-based reasoning significantly improves the effectiveness of fashion recommendation systems compared to traditional approaches.

***IndexTerms* - Fashion recommendation system, semantic search, large language model, sentence transformers, multimodal personalization, outfit generation, colour compatibility, vector database, ChromaDB, occasion detection.**

I. INTRODUCTION

INTRODUCTION

Online fashion shopping has become increasingly popular, yet many users still struggle to find relevant products that match their actual needs. Most existing e-commerce platforms rely heavily on keyword-based search, where users must enter exact product terms to retrieve results. However, in real-world scenarios, users often describe their requirements using natural language, such as “outfit for a wedding” or “dress for a festival,” which may not directly match product metadata. As a result, many relevant items are not displayed, leading to poor user experience and reduced satisfaction. To address this limitation, recent advancements in artificial intelligence have introduced techniques such as semantic search, large language models, and vision-based analysis. These technologies enable systems to understand user intent rather than relying only on exact keyword matching. However, most existing solutions apply these technologies independently and do not combine them into a unified system capable of handling real-world fashion queries effectively.This paper presents FashionAI, an AI-powered fashion recommendation system designed to understand user queries in natural language and provide personalized outfit suggestions. The system integrates multiple technologies, including a large language model for query understanding, semantic vector search for product retrieval, and a rule-based engine for outfit generation. It also supports multimodal input, allowing users to upload images for additional personalization based on attributes such as skin tone, age group, and body type. Unlike traditional systems that recommend individual products, the proposed system generates complete outfit combinations that reflect real-world dressing styles. It also incorporates cultural and seasonal awareness to better handle diverse user queries, especially those related to festivals and regional occasions. By combining semantic understanding, personalization, and outfit intelligence, the system significantly improves the relevance and usability of fashion search. The system is evaluated across multiple scenarios involving different age groups, genders, and occasions. The results demonstrate that the proposed approach provides accurate recommendations, consistent outfit generation, and fast response times. This work highlights the importance of integrating multiple AI techniques to build more intelligent and user-centric fashion recommendation systems.

II. NEED OF THE STUDY.

The rapid growth of online fashion shopping has created a problem where users are unable to find suitable outfits easily. Most existing e-commerce platforms depend on keyword-based search, which cannot understand user queries properly. Users often describe their needs in natural language such as “wedding outfit” or “beach wear,” but current systems fail to interpret these queries accurately. In addition, existing platforms do not provide complete outfit suggestions or personalized recommendations based on user attributes like age, skin tone, and body type. Users have to manually search and combine products, which is time-consuming and confusing. Therefore, there is a need for an intelligent system that can understand user queries, provide personalized outfit recommendations, and improve the overall shopping experience.

2.1 POPULATION AND SAMPLE

The dataset used in this project consists of fashion products collected from available sources. A total of 5,690 products are included in the system, covering different categories such as shirts, trousers, dresses, ethnic wear, footwear, and accessories. The products are classified based on attributes like gender (Men, Women, Boys, Girls), usage (Casual, Formal, Party, Sports, Ethnic), colour, and price. The dataset is designed to ensure that each category has sufficient number of products so that the system can generate proper recommendations and outfit combinations.

2.2 DATA AND SOURCES OF DATA

The project uses secondary data collected from fashion product datasets. The data includes product details such as product name, category, colour, gender, usage, and price.The data is processed and stored in two formats. Structured data is stored in SQLite database for fast filtering and retrieval. Semantic embeddings of products are generated using a sentence transformer model and stored in ChromaDB. This allows the system to perform both keyword-based and semantic search efficiently.

2.3 THEORETICAL FRAMEWORK

The system is based on a combination of artificial intelligence techniques including natural language processing, semantic search, and rule-based logic.User input (text or image) is processed to extract important information such as gender, occasion, and product type. Semantic embeddings are used to find similar products based on meaning rather than exact keywords. A ranking system is applied using multiple factors such as similarity score, gender match, occasion match, colour compatibility, and user preferences. Finally, a rule-based engine generates complete outfit combinations by selecting suitable items from different categories. This approach improves the accuracy of recommendations and provides a better user experience compared to traditional systems.

III. RESEARCH METHODOLOGY

The methodology section explains the design and working process of the FashionAI system. It describes how the data is collected, processed, and used to generate personalised fashion recommendations. The system combines natural language processing, semantic search, and rule-based logic to improve accuracy and user experience.

3.1 Population and Sample

The dataset used in this study consists of fashion products collected from available sources. The system includes a total of 5,690 products covering different categories such as shirts, trousers, dresses, ethnic wear, footwear, and accessories. The products are categorized based on attributes such as gender (Men, Women, Boys, Girls), usage (Casual, Formal, Party, Sports, Ethnic), colour, and price. The dataset ensures proper coverage across different combinations of product types and user categories so that the system can generate accurate recommendations and outfit combinations.

3.2 Data and Sources of Data

The project uses secondary data collected from fashion product datasets. The data includes product details such as product name, category, colour, gender, usage, and price.The data is processed and stored in two formats. Structured data is stored in a SQLite database for efficient filtering and retrieval. Semantic embeddings are generated using a sentence transformer model and stored in ChromaDB. This enables the system to perform both keyword-based search and semantic similarity search effectively.

3.3 THEORETICAL FRAMEWORK

The proposed system is based on a hybrid approach combining artificial intelligence techniques such as natural language processing, semantic search, and rule-based decision making. User queries (text or image) are processed to extract important attributes such as gender, occasion, and product type. Semantic similarity is calculated using vector embeddings to identify relevant products. A multi-signal ranking method is applied based on factors such as similarity score, occasion match, gender match, colour compatibility, and user preferences.

3.4 SYSTEM DESIGN AND IMPLEMENTATION

This section describes the design and working process of the FashionAI system. The system is developed using a multi-layer architecture that integrates artificial intelligence, semantic search, and rule-based logic to generate personalised fashion recommendations. The implementation is divided into different stages including query processing, search enhancement, product retrieval, ranking, and outfit generation. Each stage is designed to improve accuracy and provide relevant results based on user input. The system supports both text-based and image-based inputs. User queries are processed using a language model to extract important attributes such as gender, occasion, and product type. These attributes are then used in the search pipeline to retrieve suitable products.

3.4.1 Query Processing and Understanding

The first step in the system is query processing. The user input, which may be a text query or an image, is analysed to extract meaningful information. A language model is used to convert natural language input into structured parameters such as gender, occasion, colour, and price range. If the system is unable to extract complete information, it uses rule-based methods to identify missing attributes. This ensures that the system can handle both simple and complex queries effectively.

3.4.2 Search Enhancement and Feature Extraction

After query processing, the search text is enhanced using additional information such as season, cultural context, and user preferences. The system includes predefined rules to add relevant keywords based on the detected occasion or season. If the user uploads an image, the system extracts features such as skin tone, body type, and age group. These features are stored and used to personalise future recommendations within the same session.

3.4.3 Product Retrieval and Semantic Search

The system uses semantic search to retrieve relevant products. Sentence Transformers generate embeddings for both user queries and product descriptions. These embeddings are stored in ChromaDB and compared using cosine similarity. The system also applies filters such as gender, usage, and price range to narrow down the results. This combination of semantic search and filtering improves both accuracy and efficiency.

3.4.4 Ranking and Recommendation

The retrieved products are ranked using a multi-signal ranking algorithm. The ranking considers several factors including similarity score, occasion match, gender compatibility, colour matching, and user preferences. Products that match the user's requirements more closely are given higher priority. This ensures that the final recommendations are relevant and personalised.

3.4.5 Outfit Generation

The system generates complete outfit combinations using a rule-based approach. Products are classified into categories such as tops, bottoms, footwear, and outerwear. Based on the detected occasion, suitable combinations are created. Colour compatibility rules are applied to ensure visually appealing outfits. Each outfit is assigned a score based on style, completeness, and compatibility. The system returns a set of diverse outfit suggestions to the user.

IV. RESULTS AND DISCUSSION

4.1 Results of Descriptive Statics of Study Variables

Table 4.1: Descriptive Statics

Metric	Total Cases	Passed	Accuracy (%)	Remarks
Occasion Detection	48	48	100%	All cases correctly matched
Gender Mapping	5690	5690	100%	No cross-gender errors
Colour Compatibility	174 pairs	174	100%	Fully symmetric matching
Wardrobe Matching	15	15	100%	Correct pairing results
API Response Reliability	30	30	100%	No failures observed

The results shown in Table 4.1 represent the performance evaluation of the FashionAI system across multiple functional components. The system was tested using different scenarios to verify accuracy, reliability, and consistency. The occasion detection module was evaluated using 48 different combinations of age group, gender, and usage category. All test cases returned correct results, achieving 100% accuracy. This confirms that the system correctly interprets user queries and maps them to appropriate fashion categories. The gender mapping process was validated across all 5,690 products in the dataset. The system successfully ensured that no cross-gender recommendations were generated, maintaining complete consistency in results. The colour compatibility module was tested using 174 predefined colour pairs. All combinations were verified to be symmetric, ensuring that outfit suggestions maintain visual harmony. Wardrobe matching functionality was tested using 15 different clothing scenarios. In all cases, the system generated correct complementary product suggestions, demonstrating effective outfit pairing logic. API response reliability was also tested under multiple requests. All API endpoints returned valid responses without errors, confirming system stability and robustness. Overall, the results indicate that the system performs accurately and efficiently across all tested components. The high accuracy across different modules demonstrates the effectiveness of combining semantic search, rule-based logic, and AI-based processing in the FashionAI system.

4.2 Results Discussion

The results obtained from the system evaluation demonstrate that the FashionAI platform performs effectively across all major functionalities. The combination of semantic search, rule-based logic, and AI-driven processing contributes to high accuracy and reliable performance. The occasion detection module showed strong performance by correctly identifying all tested scenarios. This indicates that the system can understand both simple and complex user queries, including cultural and context-based inputs. The use of keyword mapping along with AI-based parsing improves the system's ability to interpret user intent accurately. The gender mapping process ensured that products were filtered correctly without any cross-gender mismatches. This is important in fashion recommendation systems, where incorrect categorization can lead to poor user experience. The normalization and validation logic used in the system helped maintain consistency across all results. Colour compatibility and outfit generation modules also produced highly reliable results. The use of predefined colour rules ensured that all outfit combinations were visually balanced. Unlike random product suggestions, the system generates coordinated outfits that follow real-world fashion guidelines. The ranking mechanism played a key role in improving recommendation quality. By combining multiple factors such as similarity score, occasion relevance, and user preferences, the system was able to prioritize the most relevant products. This multi-signal approach provided better results compared to single-factor ranking methods. The system also demonstrated strong reliability through consistent API responses and stable performance during testing. The fallback mechanism ensured that even in cases where AI processing fails, the system still returns meaningful results instead of empty outputs. Overall, the results confirm that the proposed system successfully addresses the limitations of traditional fashion search systems. It provides accurate, personalized, and context-aware recommendations, making it suitable for real-world e-commerce applications.

V. Acknowledgment

The authors would like to express their sincere gratitude to Mrs. M. Priyadharshini, Assistant Professor (O.G), for her valuable guidance, continuous support, and constructive suggestions throughout the development of this project. Her insights and encouragement played a significant role in completing this work successfully. The authors also thank the Department of Computer Science and Engineering, SRM Valliammai Engineering College, for providing the necessary resources and environment to carry out this project. Finally, we extend our appreciation to all those who contributed directly or indirectly to the successful completion of this work.

VI. References

[1] Reimers, N., & Gurevych, I., "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing, 2019.
[2] Wei, J., et al., "Finetuned Language Models Are Zero-Shot Learners," Proceedings of the International Conference on Learning Representations (ICLR), 2022.
[3] Alayrac, J. B., et al., "Flamingo: A Visual Language Model for Few-Shot Learning," Advances in Neural Information Processing Systems (NeurIPS), 2022.
[4] Johnson, J., Douze, M., & Jégou, H., "Billion-scale similarity search with GPUs," IEEE Transactions on Big Data, 2021.
[5] Malkov, Y. A., & Yashunin, D. A., "Efficient and robust approximate nearest neighbor search using Hierarchical Navigable Small World graphs," IEEE Transactions on Pattern Analysis and Machine Intelligence, 2020.
[6] Han, X., et al., "VITON: An Image-based Virtual Try-on Network," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018.
[7] Kang, W. C., et al., "Visually-Aware Fashion Recommendation and Design with Generative Image Models," IEEE International Conference on Data Mining, 2019.
[8] Gao, Y., et al., "Recommendation Systems with Generative Models: A Survey," ACM Transactions on Information Systems, 2023.
[9] Hershcovich, D., et al., "Challenges in Cross-Cultural Natural Language Processing," Association for Computational Linguistics (ACL), 2022.
[10] Karpukhin, V., et al., "Dense Passage Retrieval for Open-Domain Question Answering," Proceedings of EMNLP, 2020.