IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 1, January 2026

DOI: 10.17148/IJARCCE.2026.15111

LIGHTWEIGHT IMAGE BASED FASHION RECOMMENDER

Mrs. Bhagyashri K Kulkarni[1], Anagha S[2], Avinash Nayak M[3], Diya Ajith Kasbekar[4], Likitha R[5]

Department of Computer Science & Engineering,

K S School of Engineering and Management, Bengaluru, Karnataka[1-5]

Abstract: ELARA is a lightweight outfit recommendation system that combines explainable styling rules with basic image processing. It creates embeddings from clothing images using MobileNetV2 and compares them to fashion datasets using cosine similarity to determine the type of item. ELARA creates rule-based outfit recommendations based on this match to actual styling principles. Additionally, users can search datasets for suggestions. To determine appropriate and inappropriate color schemes, the system incorporates a webcam-based color analysis tool and a fashion chatbot. All things considered, ELARA provides a quick, convenient and easy-to-use solution for daily fashion advice.

Keywords: MobileNetV2, Image embeddings, Cosine similarity, Chatbot.

I. INTRODUCTION

Many people find it difficult to choose clothes that match each other in terms of colour, style and variety. An intelligent system that can help with outfit planning is becoming more essential as digitalization and online shopping increases. Traditional machine-learning techniques are less useful for real-time applications because they often require large-scale datasets and intensive computations. To solve this problem, ELARA is presented as an intelligent, lightweight outfit recommendation system that prioritizes rule-based fashion reasoning and basic visual understanding over complex deep learning models.

MobileNetV2, which is pre-trained on ImageNet, is used by ELARA as a lightweight embedding extractor. The model transforms a user-uploaded clothing image into a compact feature vector that reflects its visual characteristics. The closest matching item is found by comparing this embedding with a pre-computed dataset embedding using cosine similarity. A rules-based fashion engine is fed metadata of matched items, including product type and category, to generate complete outfit recommendations. This method guarantees that each combination recommended remains logical, consistent, and appropriate from a stylistic point of view.

ELARA includes a webcam-based color-analysis module that assesses skin tone and suggests appropriate color schemes, as well as a text-based chatbot that can provide explanations, style advice and fashion tips to further enhance user interactions. With these additional features, ELARA becomes a complete virtual assistant designed to improve the process of choosing an outfit.

II. METHODOLOGY

The ELARA smart outfit recommendation system integrates computer vision, rule-based fashion logic and interactive support, to provide a comprehensive styling platform. Dataset creation, feature extraction, similarity matching and recommendation generation are the four main stages of its pipeline. Personalization and user experience are improved by additional features like color analysis and an integrated chatbot. All things considered, the system is lightweight, easy to understand, and suitable for real-time customer-facing applications.

Product photos and metadata were scraped from the H&M online catalog to create the dataset. Due to uneven representation of categories a balanced subset of approximately 1,300 images was generated from the original large scrape using a proportional sampling technique. Duplicates, corrupted files, poor quality photos, and entries without metadata were manually removed. After cleaning, the images were resized to 224 by 224 pixels and normalized to ensure consistent feature extraction. To generate organization recommendations, ELARA uses this final curated dataset as its reference library.

Using MobileNetV2, each clothing image is converted into a 1280-dimensional feature vector that captures color, texture, shape, and pattern. For quick retrieval, all dataset embeddings are pre-computed and stored. The same process is used when a user uploads an image, and cosine similarity is used to find the closest match. This method minimizes classification errors, handles a wide variety of clothing styles, and allows real-time performance.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 1, January 2026

DOI: 10.17148/IJARCCE.2026.15111

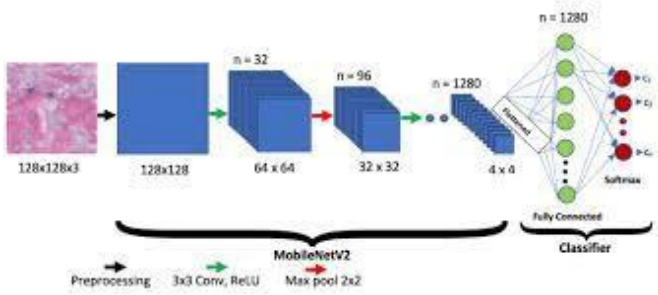

Fig. 1. MobileNetV2 Architecture

As the main feature extractor for clothing images, MobileNetV2 is essential for ELARA. Each clothing image is transformed into a compact 1280-dimensional embedding that captures texture, shape, color patterns, and overall visual style using the network's pre-trained layers, which were initially trained on ImageNet showed in Fig.1. For each item in the dataset, these embeddings act as unique "visual signatures". When a user uploads an image, it is processed by the same MobileNetV2 pipeline, which creates an embedding that is compared with all stored dataset embeddings using cosine similarity. The recommendation logic is decided by closest match, which also establishes the item category.

ELARA proceeds to its rule-based recommendation phase after retrieving the most similar dataset items. This module determines which complementary fashion items should be suggested based on the matched item's metadata, specifically its category (tops, bottoms, full-body, shoes, accessories) and product type. The guidelines emulate the logic of real-world styling: clothes are paired with bags, accessories, and shoes; The top is paired with bottoms and shoes; And accessories are matched based on their functional role. This rule-based engine ensures that recommendations are practical and stylistically consistent, prevents fashion mismatches, and ensures consistency.

ELARA includes a conversational chatbot that uses Generative AI (Gen AI) as an auxiliary intelligence layer in addition to the recommendation system. This text-based interface extends the system beyond static output, by providing natural, human-like guidance on a range of fashion-related questions. Specifically, the chatbot provides immediate assistance on a variety of topics, such as advanced color theory application, seasonal styling advice, detailed outfit coordination, and garment care instructions. The chatbot significantly improves the overall usability of ELARA by offering dynamic, on-demand support, transforming the program from a simple automated generator to a customized, interactive styling assistant. By answering complex, context-specific questions that fall outside the scope of a purely visual recommendation engine, this conversational capability increases user engagement and usefulness.

Performance Evaluation Functions

Traditional metrics such as accuracy, precision, recall or confusion matrix are not directly applicable to ELARA as it serves as a similarity-based retrieval system rather than a classification model.

Rather, the system is evaluated on the logical quality of the rule-based organization recommendations as well as the accuracy and consistency of the matches obtained. Reliable nearest-neighbor matching across different clothing categories is made possible by the MobileNetV2 embedding pipeline, which generates static 1280-dimensional feature vectors. User testing demonstrates that rule-based recommendations are consistent with sensible fashion styling logic, and qualitative inspection verifies that visually similar items are consistently retrieved. The system is suitable for interactive use in the real world because it responds in real time, with embeddings and similarity calculations taking less than a second.

III. DATASET AND PARAMETER SETTINGS

All online product photos from the H&M website were used to create the dataset for the ELARA smart outfit recommendation system. Several thousand fashion items from categories such as tops, bottoms, dresses, shoes, bags, jewelry, and accessories were retrieved during the initial web-scraping process. However, there was a significant imbalance in the raw data, with certain categories being dominant. A balanced subset of approximately 1,308 carefully selected photographs was extracted using a proportional sampling technique to guarantee similarity and avoid biased similarity matching.

Each item was manually checked before inclusion to eliminate duplicates, corrupted files, incorrect links, blurry images, and entries lacking essential metadata. Only excellent samples with comprehensive textual information including product type, garment group, color group and product description were retained. Reliability during feature extraction and rule-based recommendation was guaranteed by this cleaning step.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 1, January 2026

DOI: 10.17148/IJARCCE.2026.15111

Each image was normalized to match the input format required by MobileNetV2 and resized to a consistent 224x224 resolution. Even though the system does not rely on fully supervised classification, the curated dataset was divided into training, validation, and test sets for internal assessment of the feature extraction process and the stability of similarity retrieval.

The dataset serves as a reference library rather than a label-based training corpus as the recommendation engine relies primarily on similarity search and metadata mapping.

MobileNetV2's pre-trained feature extractor was used to calculate the embeddings for each dataset item. To facilitate quick similarity search during runtime, these 1280-dimensional vectors were stored in a compressed format.

Table 1. Category Distribution of Curated Dataset

Category	Count
Tops	~300
Bottoms	~250
Full body	~200
Footwear	~180
Accessories	~250
Jewellery	~120

Table 1 presents the cleaned clothing images that were used for our experiments in recommendation system.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

Several functional aspects of the ELARA smart outfit recommendation system were evaluated, including chatbot interactions, rule-based outfit generation, similarity-based retrieval, and real-time color analysis. The evaluation focuses on qualitative performance, response stability, and system behavior under real-world usage scenarios rather than traditional accuracy metrics because ELARA is not a classification model.

- *Extracting features and matching similarities*

MobileNetV2, which was pre-trained on ImageNet and used exclusively as a feature extractor rather than a classifier, processed all the cleaned dataset images. To capture shape, texture, color, style, and scene structure, each image was transformed into a 1280-dimensional embedding. These embeddings were created once and saved for quick access.

The embedding generation process demonstrated lightweight and efficiency during evaluation, averaging 0.05–0.1 seconds per image on a typical CPU environment. The same embedding pipeline is used by ELARA when a user uploads an image, and cosine similarity is calculated against all stored embeddings. Testing showed that similarity search consistently returned the correct or visually closest item, even when photos were rotated, captured in different lighting conditions, or had backgrounds. This method enables strong performance in overlooked fashion styles or non-standard images and removes the need for traditional classification.

- *Rule Based Recommendation Behavior*

ELARA's rule-based engine decides which complementary apparel items should be suggested after finding the closest dataset items. Experiments conducted on several outfit combinations confirmed that the rule logic behaves consistently across item categories. Examples include:

- Tops → paired with bottoms, shoes, and accessories.
- Full-body items (dresses, jumpsuits) → paired with footwear and bags.
- Bottoms → paired with tops and shoes
- Shoes → paired with full outfits rather than other shoes.
- Accessories → paired based on functional relevance. Testing across all key categories showed that the rules generated practical and stylistically consistent outfit recommendations that mimicked human stylist choices without the need for sophisticated deep learning models.

- *Color Analysis Evaluation*

A webcam-based colour analysis module in ELARA determines the user's undertone category and suggests appropriate clothing colours. Multiple-user experiments conducted in various indoor lighting conditions revealed:

- *Consistent recognition of dominant facial colors*
- *Regular classification into generalized enterprise groups*
- *Customized color recommendations according to generally recognized color theory*

DOI: 10.17148/IJARCCE.2026.15111

The module effectively offered practical, real-time color guidance that, despite not being designed for professional colorimetry, improved the overall fashion decision-making experience.

- *Chatbot Interaction and User Feedback*

The responsiveness and relevance of a specially designed chatbot with a fashion focus incorporated into the interface was evaluated. Questions about styling advice, outfit pairing, color matching, clothing care and seasonal trends can be answered by the chatbot.

Strong user satisfaction was demonstrated by user feedback, which described the chatbot as "easy to use," "helpful" and "similar to browsing advice on fashion apps." The conversational style of the chatbot enhances the conversation to a great extent and reduces the user's confusion while choosing an outfit.

- *System usability, speed and performance*

The system response was evaluated during several full-pipeline tests, from image upload to final outfit recommendation. The results were consistent:

- Embedding Extraction Time: ~0.1–0.2 seconds
- Similarity Search Time: <0.05 seconds
- Rule-Based Recommendation Time: <0.01 seconds
- Total Response Time: 0.2–0.4 seconds on average
- Interface Feedback: Users found ELARA intuitive, visually appealing, and aligned with typical online shopping workflows.

These findings show that ELARA can work without GPU acceleration and is suitable for real-time applications.

- *Limitations*

Despite the good performance of ELARA, several limitations were found:

- Due to the small size of the dataset (approximately 1300 items), diversity in specific fashion categories is limited.
- The system only suggests items from curated datasets; It is unable to make original costumes.
- If user photos with too much background are not manually cropped, the similarity accuracy may be reduced.
- Color analysis may change depending on lighting.

These restrictions provide opportunities for further development.

- *Overall Findings*

The experimental evaluation shows that ELARA effectively creates an effective styling assistant by fusing computer vision, fashion logic, and user-centric interaction. The rule-based engine regularly generated logical and aesthetically pleasing outfit combinations, and the MobileNetV2 embedding pipeline demonstrated strength and portability. ELARA is now a useful and entertaining smart outfit recommendation system thanks to the chatbot and colour analysis modules, which further improved usability.

ELARA provides dependable real-time performance despite dataset limitations and offers a promising basis for more sophisticated, scalable fashion-tech applications.

V. CONCLUSION

The ELARA smart outfit recommendation system effectively integrates computer vision, rule-based fashion reasoning, and interactive user support to improve the digital fashion-browsing experience. Using MobileNetV2 as a lightweight feature extractor, the system efficiently converts clothing images into high-dimensional embeddings, enabling accurate similarity matching without the need for a full classification model. A carefully selected dataset of 1,300 fashion items serves as a trusted reference library from which recommendations are generated in real-time.

The system is further enhanced by a rules-based engine, which guarantees that outfit recommendations are stylistically consistent and represent useful, real-world fashion standards. The color analysis module and fashion-focused chatbot are two other features that elevate ELARA from a basic recommendation tool to a full-featured virtual styling assistant.

Overall, the experimental findings show that ELARA provides quick responses, recommendations that are visually consistent, and offers an easy-to-use interface. The results demonstrate that despite the system's limitations due to the size of the dataset and the lack of deep learning classification metrics, embedding-based retrieval in combination with rule-driven logic provides a reliable and understandable method for personalized fashion assistance. Larger datasets, more fashion categories, seasonal recommendations, and improved generator styling capabilities could be investigated in future work.

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 1, January 2026

DOI: 10.17148/IJARCCE.2026.15111

REFERENCES

[1]. S. Shirkhani, H. Mokayed, R. Saini, and H. Y. Chai, "Study of AI-driven fashion recommender systems," *SN Computer Science*, vol. 4, no. 5, p. 514, 2023.

[2]. D. Onitiu, "Algorithmic abstractions of 'fashion identity' and the role of privacy with regard to algorithmic personalisation systems in the fashion domain," *AI & Society*, vol. 37, no. 4, pp. 1749–1758, Dec. 2022.

[3]. S. Wazarkar, V. Chavan, and U. Gokhale, "Advanced fashion recommendation system for different body types using deep learning models," *International Journal of Engineering Research & Technology (IJERT)*, vol. 11, no. 4, pp. 1–5, 2022.

[4]. S. Chakraborty, R. Roy, S. Manna, and S. Basu, "Fashion recommendation systems, models and methods: A review," *Informatics*, vol. 8, no. 3, p. 54, 2021.

[5]. R. Sulthana, "A review on the literature of fashion recommender system using deep learning," *International Journal of Performability Engineering*, vol. 17, no. 8, pp. 695–703, 2021.

