iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: V Month of publication: May 2026

DOI: https://doi.org/10.22214/ijraset.2026.83066

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 V May 2026- Available at www.ijraset.com

AI-Powered Smart Wardrobe and Outfit Recommendation System

Keshav Sonawane[1], Atharva Kotkar[2], Satyam Bartakke[3], Rugved Shinde[4], Dr. Vaishali Suryawanshi[5]
Department of Computer Engineering and Technology Dr. Vishwanath Karad MIT World Peace University Pune, India

Abstract: Ever stared at your closet and thought, "Why do I keep reaching for the same old shirt?" Yeah, you're not alone. Managing a growing wardrobe turns into this messy daily scramble, and honestly, it gets annoying fast. That's why we built an AI-powered Smart Wardrobe and Outfit Recommendation System—to help you keep track of your clothes and make picking out an outfit feel less like a chore. Here's how it actually works: Just snap a picture of your clothes and upload them. Our system sorts everything into 15 categories using the ConvNeXt-Large deep learning model. We picked this one because, during early tests, it handled tricky lookalikes way better than other models—like telling apart casual shirts from polos, which most systems can't seem to figure out. Instead of spitting out generic fashion advice, the system leans on a lightweight neural compatibility model. It looks at real, curated examples to learn what pieces fit together naturally. So when it suggests an outfit, it's not just random—it's actually something you'd want to wear. The technical setup is pretty simple. The React frontend lets you manage your digital closet, and the Flask backend does all the image crunching and AI calculations. The model isn't perfect—it can get confused when your clothes are nearly identical—but overall, it nails classification and serves up outfit ideas people genuinely like. At the end of the day, this isn't just about making life easier. We're showing how AI can take the stress out of getting dressed and laying the groundwork for even smarter, more personal fashion tools in the future.

Index Terms: Fashion Recommendation, Clothing Classification, ConvNeXt, Outfit Compatibility, Computer Vision, Deep Learning, Smart Wardrobe System, Flask Backend, React Frontend.

I. INTRODUCTION

Many visually demanding tasks, including object recognition, image classification, and pattern analysis, can now be practically automated thanks to recent developments in deep learning and computer vision. Plenty of fashion platforms now use advanced methods for things like classifying clothes, spotting attributes, and recommending outfits. Still, lots of systems stick to old-school collaborative filtering or just go by rules. So, most of their recommendations end up relying on past user actions without really looking at the visual features of the clothing itself. Because of this, these systems frequently find it difficult to manage a variety of personal wardrobes or generate visually appealing outfit recommendations from user-uploaded photos.

The growing availability of massive, annotated datasets like DeepFashion [2] and Polyvore [3] has completely reshaped fashion analysis, pushing the field toward data-driven methods. With this wealth of information, deep learning models can now figure out visual styles and item compatibility straight from the images themselves. At the heart of this transition are Convolutional Neural Networks (CNNs), which have become essential because of how well they extract detailed spatial features from clothing photos. ConvNeXt [9] and EfficientNet [10] are two examples of modern CNN architectures that have demonstrated strong performance in detailed clothing classification tasks, even under different lighting, poses, and texture variations.

Despite these developments, there are a number of real-world obstacles when applying AI methods to personal wardrobe management. Curated catalog images and user-uploaded images are frequently very different. Partial occlu-sions, uneven lighting, and cluttered backgrounds are a few examples. In real-world situations, these factors make accurate multi-class clothing classification more difficult. Furthermore, limited or manually curated explicit compatibility annotations make it challenging to generate meaningful outfit recom-mendations. Initial tests revealed that multimodal embedding models such as CLIP [11] do not always capture the garment-level compatibility patterns required for personal wardrobe situations, despite the fact that these models offer strong semantic alignment between text and images. Instead of focusing on a large-scale retail setting, we focus on a functional and user-friendly wardrobe setting in this work. For personal wardrobe management, the suggested system integrates outfit recommendation and clothing classification into a single pipeline. A balanced 15-category Kaggle fashion dataset [17] is used to train a ConvNeXt-Large model, which uses extensive data augmentation and EfficientNet-based pre-processing to improve robustness. Furthermore, based on pre-dicted garment categories, a lightweight neural compatibility model is presented to learn top-bottom pairing relationships and produce ranked outfit recommendations.

The proposed system aims to provide contextually relevant outfit recommendations and accurate clothing classification by combining a compact compatibility learning module with a high-capacity visual feature extractor. Ultimately, we designed this framework to be highly scalable, making it a natural fit for real-world applications such as digital wardrobe management, AI-driven outfit planning, and personalized styling tools.

II. RELATED WORK

Fashion image analysis and outfit recommendation have really taken off in things like online shopping, virtual styling, and personalized shopping apps. In the beginning, researchers relied on hand-crafted features—so, stuff like color his-tograms, texture details, and simple outline shapes. They'd plug these into classic machine learning models like SVMs or k-NN to classify clothes or help people find matching items [1]. These early methods did okay in tidy lab setups, but once they hit real-world images—where people strike odd poses, backgrounds get messy, and lighting changes all over the place—they just didn't keep up.

Deep learning in fashion analysis really took off when datasets like DeepFashion [2] and Polyvore Outfits [3] showed up. Suddenly, researchers had access to tons of images, all labeled and ready to go. That's when Convolutional Neural Networks (CNNs) started dominating the field. Unlike older methods that needed a lot of manual feature selection, CNNs just dive in and learn the important visual clues directly from the images. Over the years, everything from VGG and ResNet to EfficientNet and ConvNeXt [9], [10] has pushed the field forward, getting better and better at spotting the tiny details—like textures, shapes, and patterns—that really matter for sorting out different kinds of clothing. These days, neural networks like these are pretty much the heart of fashion recognition because they just keep outperforming the old-school approaches, hands down.

Outfit recommendation systems have come a long way, too. They used to be clunky—just a bunch of hard-coded rules saying which colors or styles had to go together, so there wasn't much room for personal taste [4]. Now, compatibility models actually learn from data, and they're much better at making smart suggestions that feel tailored to what people actually want. Neural networks trained on curated fashion datasets and visual embeddings are used in recent methods to learn compatibility relationships between garments, partic-ularly for top and bottom pairings [3]. While some studies use multimodal architectures such as CLIP to link textual descriptions with visual features, others have investigated graph-based representations or attention mechanisms to model style coherence [11].

Despite these advancements, many limitations still exist in current approaches. Many methods rely heavily on large anno-tated compatibility datasets or multimodal training resources, which can be impractical for personal wardrobe scenarios. Additionally, most previous work treats clothing classification and outfit recommendation as separate issues, creating systems that lack seamless integration. Furthermore, models trained on curated catalog images may not perform consistently with user-uploaded images, which tend to show more variability.

On the other hand, this work integrates outfit recommen-dation and clothing classification into a single framework intended for personal wardrobe management. The proposed system learns garment categories and pairing relationships in a single cohesive process by combining a lightweight neural compatibility model with a ConvNeXt-Large classifier trained on a balanced Kaggle fashion dataset [17]. While maintaining high performance, this approach lessens dependence on large fashion graphs and multimodal corpora. Because of this, it's a perfect fit for real-world digital closets and personalized styling apps.

III. METHODOLOGY

We built this system to bring together image preprocessing, deep learning, and smart outfit recommendations—all in one practical pipeline. The idea? Make an AI wardrobe assistant that checks out photos you upload, figures out what each clothing item is, and then suggests outfit combos that actu-ally look good. We didn't want to mess with confusing or bloated, data-heavy methods. Instead, we kept things solid, simple, and easy to use—so it's not just for research labs but for real people's closets. We leaned on solid research in fashion recognition and outfit matching, taking inspiration from projects like DeepFashion [2], Polyvore outfit modeling [3], and newer convolutional models like ConvNeXt [9]. The whole process comes together in a pretty straightforward way: get the dataset ready, prep the images, classify the clothes, suggest outfits, and tie the system together in the end.

A. Dataset and Preprocessing

For the clothing classification, we used a publicly available Kaggle dataset that breaks garments down into 15 distinct categories [17]. It contains exactly 500 images per class, giving us a perfectly balanced dataset of 7,500 RGB images in total. Because the data is so evenly split, it ensures that everyday

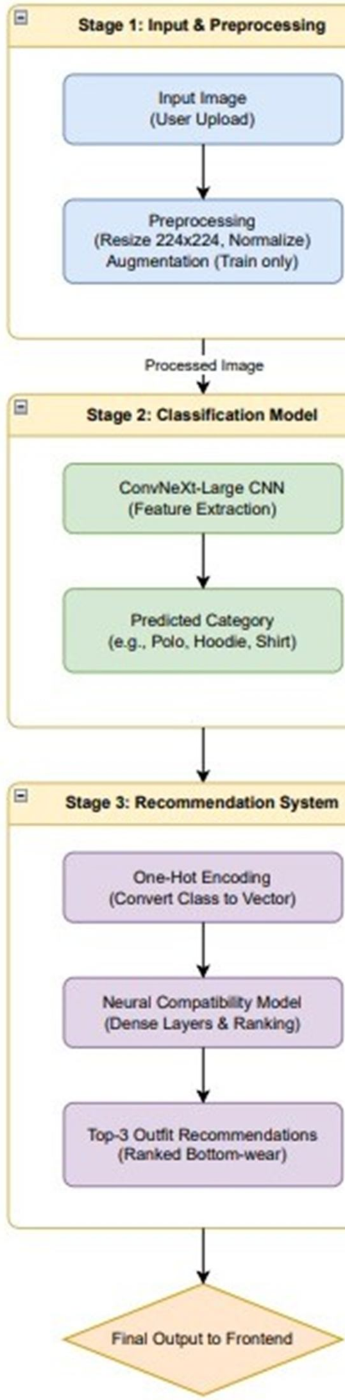

Fig. 1 Overview of the proposed end-to-end workflow from image upload to outfit recommendation.

items—like shirts, polos, hoodies, jackets, denim, jeans, skirts, and dresses—all get a fair and equal representation during training. The dataset reflects the diversity found in real-world wardrobe photos by displaying discernible variations in lighting, colours, textures, and fabric types.

We used stratified sampling to maintain class balance and divided the dataset using an 85% to 15% train-validation ratio in order to prepare the data for training. To meet the ConvNeXt input requirements and enable effective batch processing, all images were resized to 224 ×224 pixels. For pixel normalisation, we adhered to the EfficientNet preprocessing standard [10]. This decision was made after preliminary experiments showed that it produced more stable convergence than conventional normalisation techniques.

We used a variety of augmentation methods on the training set, including rotation, zoom, shear, brightness adjustment, and horizontal flipping. Early experiments showed that models trained without augmentation started to overfit, particularly in categories of clothing that were visually similar. To preserve their original distribution, the validation images were only resized and normalised.

TABLE I
Dataset Characteristics

Characteristic	Description
Source	Kaggle Clothing Dataset [17]
Total Images	7,500
Classes	15
Images per Class	500 each
Image Type	RGB Fashion Images
Train Split	6,375 images
Validation Split	1,125 images
Resolution	224×224 pixels

B. Clothing Classification Model

The ConvNeXt-Large architecture, a contemporary convo-lutional network, is used in the clothing classification com-ponent. By incorporating improvements motivated by Vision Transformers [9], it expands upon conventional CNNs. Ini-tially, a number of convolutional backbones were evaluated. ConvNeXt-Large demonstrated less confusion between visu-ally similar clothing, such as shirts and polos, and more stable training. Its selection for thorough clothing classification in a personal wardrobe setting was influenced by these findings. The model was set up with ImageNet-pretrained weights to take advantage of general visual features. A custom classifi-cation head was added to adjust the pretrained backbone for the clothing classification task. This head includes a Global Average Pooling layer, followed by Batch Normalization to stabilize feature distributions and Dropout to reduce overfit-ting. Fully connected layers with ReLU activation precede the final softmax layer, which outputs probability scores for 15 clothing categories. Training took place in two stages. In the first stage, the Con-vNeXt backbone was frozen, and only the classification head was trained. This approach helped the model learn clothing-specific patterns without upsetting the pretrained weights. In the second stage, all layers were unfrozen and fine-tuned with a lower learning rate. Full fine-tuning improved class separation but needed careful management of the learning rate to prevent overfitting. Early stopping, learning rate scheduling, and model checkpointing were used to maintain training stability. Model performance was assessed using Top-1, Top-3, and Top-5 accuracy metrics, aligning with evaluation practices found in fashion recognition studies [2], [6].

TABLE II
Training Configuration and Hyperparameters

Parameter	Value
Base Model	ConvNeXt-Large
Pretraining Dataset	ImageNet
Epochs	25
Batch Size	32
Optimizer	Adam
Initial Learning Rate	0.0005
Fine-Tuning Learning Rate	0.0001
Loss Function	Categorical Cross-Entropy
Callbacks	EarlyStopping, LR Scheduler, Checkpoint
Metrics	Top-1, Top-3, Top-5 Accuracy

C. Outfit Recommendation Model

In addition to clothing classification, the system includes a lightweight neural compatibility model that suggests suitable bottom-wear items for a given top-wear garment. The model takes a one-hot encoded representation of the predicted top-wear category as input and generates softmax-normalized probability scores for available bottom-wear categories.

This design keeps it simple while helping the model learn compat-ibility relationships effectively.

The suggested model learns pairing relationships directly from carefully chosen top-bottom combinations, in contrast to conventional rule-based recommendation systems [4]. The network discovered typical fashion trends during training, like wearing shirts with pants or polo shirts with jeans. Early training iterations had some incorrect pairings, but these errors were lessened by increasing training epochs and adjusting learning rates. Better generalization is possible with this learning-based approach than with fixed heuristic rules.

D. End-to-End System Integration

To handle image uploads, preprocessing, clothing classifica-tion, and recommendation inference, the entire system makes use of a Flask backend. When a user uploads an image, the ConvNeXt classifier receives the processed image after the backend completes the preprocessing pipeline. The recom-mendation module receives the predicted garment category and generates a ranked list of bottom-wear options that are visually compatible. We built the interactive frontend with React, so users can manage their wardrobe and check out outfit suggestions instantly.

This modular setup keeps everything flexible, so scaling up or adding new features down the road won't be a headache. Say we want to bring in multimodal models like CLIP [11] or layer in more contextual info — we can do that without tearing apart the whole architecture.

IV. RESULTS AND ANALYSIS

Let's dig into the experimental results for our AI wardrobe system. We looked at how the ConvNeXt-Large classifier actually did, watched its training progress, checked how well the neural compatibility model works, and made sure the whole system runs smoothly from start to finish. For testing, we set up everything on a Tesla T4 GPU, using TensorFlow and Keras [14].

A. Experimental Setup

For our 15-class clothing classification, we went with the ConvNeXt-Large model and kicked things off using weights pretrained on ImageNet [9]. We broke up the training into two parts. First, we locked down the base convolutional layers and focused on training our custom classification head for 15 epochs. With this configuration, the model was able to learn characteristics unique to clothing without interfering with the pretrained representations. All layers were unfrozen and refined for an additional ten epochs at a reduced learning rate during the second phase. To enhance generalization, extensive data augmentation and EfficientNet-style preprocessing were used during training.

B. Classification Performance

The classifier showed stable convergence during both train-ing phases. Validation accuracy rose quickly in the initial frozen phase and then gradually leveled off during fine-tuning, hitting a peak validation accuracy of **87.29%**. There were slight fluctuations near convergence, mainly due to visually similar garment categories like shirts and polos. The consis-tently high Top-3 and Top-5 accuracy values show that the model could accurately rank the true class, even when the top prediction was unclear.

TABLE III
ConvNeXt-Large Clothing Classification Performance

Metric	Score
Best Validation Accuracy	87.29%
Final Validation Accuracy	86.93%
Training Accuracy	86.60%
Top-3 Accuracy	98.40%
Top-5 Accuracy	99.56%

These results are consistent with performance trends re-ported in large-scale fashion datasets such as DeepFashion [2] and highlight the effectiveness of ConvNeXt-Large for fine-grained clothing categorization.

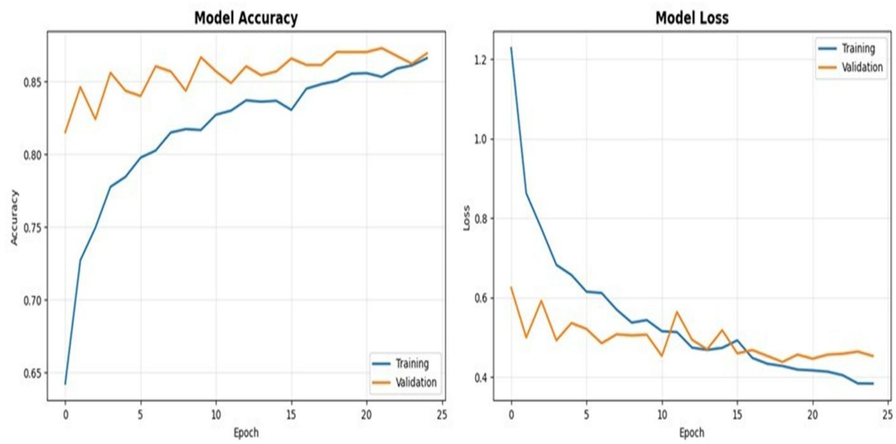

Fig. 2. Training and validation accuracy and loss curves for the ConvNeXt-Large classifier.

1) *Learning Curve Behavior:* The training and validation accuracy and loss curves, shown in Fig. 2, exhibit smooth

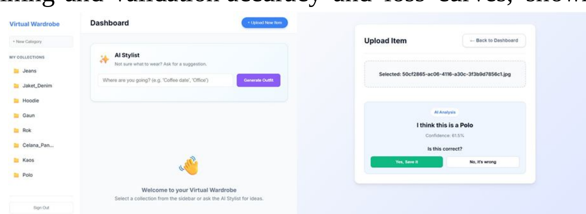

(a) Uploading a clothing image (b) Predicted top-wear category

Fig. 3 Initial stages of the smart wardrobe interface.

convergence without signs of severe overfitting. Validation per-formance plateaued during the fine-tuning stage, and the best-performing model checkpoint was obtained at epoch 22. Visual inspection of the curves suggests that the applied augmentation and regularization strategies contributed positively to training stability.

C. *Outfit Recommendation Model Results*

The lightweight neural compatibility model was trained to match predicted top-wear categories with compatible bottom-wear options. The model reached a highest training accuracy of 83.33% and showed a steady drop in loss over 100 train-ing epochs. The learned compatibility patterns aligned with commonly accepted fashion pairings, similar to findings from earlier studies on compatibility learning [3], [7]. In the early training stages, there were occasional mis-matches for less common pairings. These errors decreased after we adjusted the learning rate and increased the number of training epochs. This showed that the model improved from gradual convergence instead of rapid optimization.

D. *End-to-End Evaluation*

To evaluate the complete system, an end-to-end inference test was performed using a user-uploaded Polo shirt image. The ConvNeXt classifier predicted the garment category as *Polo* with **100% confidence**. This predicted label was then passed to the outfit recommendation model, which generated a ranked list of compatible bottom-wear items as follows:

- Jeans: 0.62
- Celana_Pendek: 0.14
- Celana_Panjang: 0.13

Based on these probabilities, the final outfit selected by the system was:

Polo (Top) + Jeans (Bottom)

This recommendation aligns with intuitive fashion prefer-ences and is consistent with results reported in similar AI-based styling systems such as AI-Fashionista [8].

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 V May 2026- Available at www.ijraset.com

E. System Demonstration

The screenshot above present a collection of user interface. These pictures show the entire process, from uploading photos

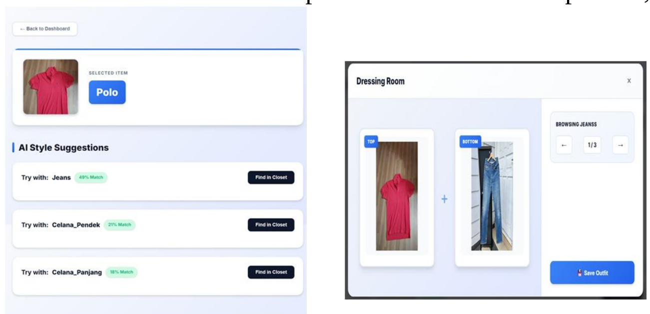

(a) Recommended bottom-wear items (b) Final assembled outfit

Fig. 4 Outfit recommendation and final selection interface.

and classifying clothing to creating outfit suggestions and putting the finished ensemble together. The outcomes verify that the system operates in real time and offers user-friendly outputs.

V. LIMITATIONS

The suggested AI-powered Smart Wardrobe and Outfit Recommendation System has a number of drawbacks, despite its high classification accuracy and insightful outfit recommendations.

First, 7,500 carefully chosen fashion photos from Kaggle make up the dataset used for training and assessment. This dataset does not accurately represent the variety of wardrobes found in real life, despite covering 15 different categories of clothing. Different body shapes, messy backgrounds, uneven lighting, layered outfits, and cultural clothing styles are exam-ples of variations that are under-represented. Because of this, the model might not work well with images created by users. Second, the system primarily assesses clothing according to its category labels and outward appearance. At the moment, it does not account for contextual or individual factors like user style preferences, body measurements, weather, occasion type, or colour harmony rules. These elements have a signif-icant impact on outfit choice in real-world styling scenarios.

Without them, recommendations might be less relevant and personalised.

Third, a comparatively small collection of carefully cho-sen top-bottom pairing examples is used to train the outfit compatibility model. The model's capacity to learn intricate or unusual fashion relationships is constrained by the small number of pairings when compared to larger fashion com-patibility datasets like Polyvore or DeepFashion. Conservative recommendations favouring common combinations may result from this.

Lastly, the system does not take into account three-dimensional body shape information, fabric draping, or gar-ment fit; instead, it relies solely on 2D image features. These elements are essential for the comfort and compatibility of clothing in the real world. These problems would require the integration of body modelling or 3D garment analysis techniques, which are outside the current purview.

VI. ETHICAL CONSIDERATIONS

AI-driven fashion analysis and recommendation systems raise a number of ethical concerns that must be addressed to guarantee responsible and reliable use. User privacy and data security have to come first. A smart wardrobe system works with personal photos, and let's face it, those images can reveal a lot about someone's life—how they dress, where they live, and what their daily routines look like. So we really need to handle that data with care. If this system actually hits the market, things like secure image storage, tight access controls, and solid anonymization aren't just nice-to-haves; they're essential. Following rules like GDPR isn't some box-ticking exercise. It matters because it genuinely protects people from having their data misused or exposed.

Let's be honest—algorithmic bias and inclusivity are huge hurdles here. Most fashion datasets stick to a pretty narrow view: they're packed with certain body types, Western styles, and what's considered "mainstream." So if your personal style or cultural background doesn't show up in the training data, these models end up missing the mark. There's another issue that's hard to ignore: the moral side. AI-powered styling tools can easily nudge people into the fast-fashion cycle, making them feel like they've always got to chase new trends and buy more stuff. That gets ugly for the planet and the wallet. Developers have to make sustainability part of the foundation—not just a side note—by encouraging people to make the most of what they already own, shop with intention, and choose clothes that last.

And, honestly, there's a danger in getting too comfort-able with letting the algorithm decide everything. Sure, an AI wardrobe assistant can help organize and inspire, but it shouldn't smother someone's unique style or creativity. If you want people to actually trust and use the system, you've got to keep it transparent about why it's suggesting certain looks, with the user always holding the reins when it comes to making the final call.

VII. CONCLUSION

In this work, we've developed a complete AI-powered wardrobe assistant that bridges the gap between deep learning classification and neural outfit recommendations. By leverag-ing the ConvNeXt-Large architecture for feature extraction, our system reached a peak validation accuracy of 87.29

Beyond just identifying clothes, our compatibility model successfully picked up on the nuances of top-and-bottom pairings by learning directly from curated examples. This allows the system to generate suggestions that actually make sense for a user's closet. With a React frontend and a Flask backend working in tandem, the entire pipeline—from photo upload to the final outfit recommendation—is smooth and responsive. While our results are promising, we did notice that performance can dip when faced with tricky lighting or items that look nearly identical. These edge cases really highlight why future iterations will need even more diverse datasets and a deeper sense of contextual awareness. Ultimately, our framework shows that pairing modern convolutional models with lightweight recommendation engines is a highly scalable way to help people organize their wardrobes and discover new styles.

VIII. FUTURE WORK

While our current results are promising, there is still plenty of room to grow. One obvious next step is expanding our training data to include a much wider variety of styles, cultural garments, and "in-the-wild" user photos. Tapping into even larger datasets like DeepFashion2 or Polyvore Outfits would definitely help the model pick up on more subtle visual cues and complex pairing logic.

We're also working on making the recommendations feel more personal. Imagine the system considering your favorite colors, your body measurements, even your go-to outfits—then suggesting things that actually feel like "you." And if we bring in context—like what the weather's doing, if it's morning or night, or what event you're dressing for—the AI suddenly gets way more practical for everyday use. On the technical side, multimodal learning opens up a ton of possibilities. By integrating models like CLIP, you could just type, "I need something formal for a winter wedding," and let the system find it for you. Long-term, we're really curious about things like 3D garment modeling and body shape estimation. If we can figure out how fabrics drape and fit a specific person, that's a whole new level for digital styling. And getting this onto a mobile app, with smooth on-device processing, would make it so much easier for people to use—like, you wake up, stand in front of your closet, and get instant, tailored ideas.

REFERENCES

[1] S. Jagadeesh, R. Piramuthu, and A. Bhardwaj, "Large-scale visual recommendations from street fashion images," in Proc. ACM Int. Conf. on Knowledge Discovery and Data Mining (KDD), pp. 1925–1934, 2014.
[2] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang, "DeepFashion: Powering robust clothes recognition and retrieval with rich annotations," in Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), pp. 1096–1104, 2016.
[3] S. Han, H. Lee, and T. Park, "Learning fashion compatibility with bidirectional LSTMs," in Proc. ACM Multimedia, pp. 1078–1086, 2017.
[4] K. Tangseng, T. Okatani, and K. Yanai, "Recommending outfits from personal closet," IEEE Access, vol. 7, pp. 89285–89293, 2019.
[5] H. S. Cho, W. S. Lee, and S. J. Lee, "Machine learning models with optimization for clothing recommendation from personal wardrobe," Sensors, vol. 21, no. 14, p. 4607, 2021.
[6] J. Choi, E. Kim, and J. Lee, "OutfitX: A deep learning framework for personalized outfit recommendations," Applied Sciences, vol. 11, no. 19,
[7] p. 9053, 2021.
[8] A. Sharma and M. Singh, "PFRS: Personalized fashion recommenda-tion system using EfficientNet and multimodal learning," International Journal of Advanced Computer Science, vol. 12, no. 4, pp. 221–231, 2022.
[9] Z. Hsiao and Y. Sun, "AI-based fashion stylist recommendation system," Journal of Fashion Technology & Textile Engineering, vol. 10, no. 2, pp. 1–8, 2022.
[10] Z. Liu, et al., "A ConvNet for the 2020s," in Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2022.
[11] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for con-volutional neural networks," in Proc. Int. Conf. on Machine Learning (ICML), pp. 6105–6114, 2019.
[12] A. Radford, et al., "Learning transferable visual models from natural language supervision," in Proc. ICML, pp. 8748–8763, 2021.
[13] M. H. Kiapour, et al., "Where to buy it: Matching street clothing photos in online shops," in Proc. ICCV, pp. 3343–3351, 2015.
[14] M. Abadi, et al., "TensorFlow: Large-scale machine learning on hetero-geneous distributed systems," arXiv:1603.04467, 2016.
[15] F. Chollet, "Keras," GitHub Repository, 2015. [Online]. Available: https://keras.io
[16] A. Paszke, et al., "PyTorch: An imperative style, high-performance deep learning library," in Advances in Neural Information Processing Systems, 2019.
[17] G. Bradski, "The OpenCV Library," Dr. Dobb's Journal of Software Tools, 2000.
[18] Kaggle, "Clothes Dataset – 15-category Fashion Image Dataset," 2023. [Online]. Available: https://www.kaggle.com/