

Di fashion: utilizing diffusion models for personalized and high-fidelity generative outfit recommendations

Krunal Deshmukh ^{1*}, Arya Khade ¹, Vaishnavi Pawar ¹, Sanika Vyawahare ¹, Ms. Arti Singh ², Mrs. Priyanka Abhale ²

¹ Dr. D. Y. Patil Institute of Engineering, Management and Research, Akurdi, Pune 44

² Dr. D. Y. Patil International University, Akurdi, Pune-44

*Corresponding author E-mail: dkunal665@gmail.com

Abstract

Artificial intelligence is being used more and more by the fashion industry to boost client interaction and customization. The use of diffusion models, a subclass of generative models, in offering high-fidelity and customized clothing suggestions is examined in this research. Diffusion models gradually convert random noise into comprehensible visual outputs, making them ideal for producing intricate, detailed pictures. They are perfect for producing realistic and eye-catching costume ideas because of their capacity to capture complex materials, patterns, and clothing styles.

We analyze how these models perform better than conventional techniques like variational autoencoders (VAEs) and generative adversarial networks (GANs) in terms of visual quality, variety, and control over the generating process. The paper also emphasizes how diffusion models may be used to create personalized clothing alternatives that represent personal preferences by taking into account user preferences including body type, style, and previous fashion decisions. We also talk about how DiFashion systems may affect many areas of the fashion business, such virtual try-ons and eco-friendly clothes. The limitations and future prospects of using diffusion models for fashion recommendation systems are identified in the paper's conclusion, with special attention paid to scalability, user involvement, and computing needs.

Keywords: Generative Outfit Recommendation; Diffusion Models; Personalized Fashion; AI-Generated Content; Fashion Image Generation; Outfit Compatibility; Difashion Model.

1. Introduction

Fashion and technology have come together to generate new opportunities for innovation and customisation, which has altered how consumers engage with clothing and style. Artificial intelligence (AI) has emerged as a key player in this change, delivering powerful capabilities for automating design, personalizing recommendations, and enhancing customer experiences using machine learning models. The use of diffusion models, which have demonstrated remarkable abilities in creating detailed, excellent images, is among the latest advancements in AI-driven fashion. Diffusion models, a subclass of generative models, steadily transform random noise into coherent data, offering an unmatched level of control and precision in the images they produce.

We aim to provide a comprehensive overview of the potential of diffusion models to revolutionize the fashion industry by analyzing their benefits and drawbacks, ranging from AI-powered sustainable fashion practices to e-commerce platforms that enable virtual try-ons. This study explores the use of diffusion models to produce high-fidelity, tailored outfit recommendations—a concept we call DiFashion. We explore how diffusion models may be used to create visually realistic ensembles and customize them for each user according to their body type, fashion history, and preferred styles.

The relative benefits of diffusion models over other generating techniques, their technological underpinnings, and their potential integration into consumer and designer fashion recommendation systems will all be covered in the sections that follow. We'll talk about the difficulties in scaling and using these technologies in various fashion companies as well as the wider ramifications of generative models for the future of fashion, such as accessibility, sustainability, and personalization.

2. System architecture




Fig. 1: Diffusion Model Architecture.

3. Overview of domain

Generative Outfit Recommendation (GOR), the project's primary objective, uses AI to create personalized and visually appealing fashion ensembles. The proposed model, DiFashion, makes use of diffusion models to create multiple fashion items simultaneously while ensuring that they are complementary to each other and meet consumer preferences.

3.1. Diffusion models for fashion

Diffusion models excel at creating intricate, high-resolution images by progressively converting noise into finely detailed pictures. Fashion designers must capture delicate textures, patterns, and materials like silk, denim, or lace to produce realistic and striking costume designs. Diffusion models' ability to provide photorealistic results gives them a clear edge in the fashion industry, where visual correctness is essential.

One major benefit of diffusion models is their systematic approach to picture production. By starting with noise and repeatedly refining it, these models offer more control over the generating process than other models, including GANs (generating Adversarial Networks). This allows for greater flexibility in ensuring that the clothes manufactured meets specific requirements, such as reflecting personal preferences or current fashion trends.

One major benefit of diffusion models is their systematic approach to picture production. By starting with noise and repeatedly refining it, these models offer more control over the generating process than other models, including GANs (generating Adversarial Networks). This allows for greater flexibility in ensuring that the manufactured apparel meets specific requirements, such as adhering to current fashion trends or taking personal preferences into account.

3.1.1. Personalized outfit recommendation

Understanding each user's unique style is the first step toward customizing. Favorite colors, patterns, styles, and even fashion labels may be mentioned. For example, the model can focus on designing outfits that suit a user's taste for minimalist styles while coming up with more daring and vibrant concepts for a customer who prefers vibrant, varied styles.

User buying patterns, including past purchases or often seen items, give the model important information about the user's preferences and assist in creating ensembles that fit their shopping habits. Diffusion models can continuously improve and update recommendations as user preferences change by learning from previous choices.

Diffusion models can utilize information about a user's body measurements, like height, weight, and proportions, to recommend clothes that are better suited to their specific body type. Because different body shapes react better to specific cuts, styles, and fits, this is significant in the fashion industry. By tailoring recommendations according to the user's body type, the algorithm raises the likelihood that the user will be happy with the recommended clothing.

Diffusion models can facilitate virtual try-ons for a more realistic experience by producing images that replicate how apparel would appear on the user's body. By utilizing this to examine how clothing might fit and hang on their body, users may be able to make better decisions about size and style.

3.1.2. Personalized outfit recommendation

Understanding each user's unique style is the first step toward customizing. Favorite colors, patterns, styles, and even fashion labels may be mentioned. For example, the model can focus on designing outfits that suit a user's taste for minimalist styles while coming up with more daring and vibrant concepts for a customer who prefers vibrant, varied styles.

User buying patterns, including past purchases or often seen items, give the model important information about the user's preferences and assist in creating ensembles that fit their shopping habits. Diffusion models can continuously improve and update recommendations as user preferences change by learning from previous choices.

Diffusion models can utilize information about a user's body measurements, like height, weight, and proportions, to recommend clothes that are better suited to their specific body type. Because different body shapes react better to specific cuts, styles, and fits, this is significant in the fashion industry. By tailoring recommendations according to the user's body type, the algorithm raises the likelihood that the user will be happy with the recommended clothing.

Diffusion models can facilitate virtual try-ons for a more realistic experience by producing images that replicate how apparel would appear on the user's body. By utilizing this to examine how clothing might fit and hang on their body, users may be able to make better decisions about size and style.

3.1.3. Diffusion models ensure high fidelity

Diffusion models are used to repeatedly convert a random noise distribution into a cohesive picture. This incremental denoising technique allows the model to enhance the output at each step by beginning with a basic structure and progressively adding more detail. This meticulous procedure ensures that each iteration improves the authenticity of the produced image, producing realistic fashion designs with pronounced curves, smooth textures, and patterns.

Diffusion models' iterative refinement process makes them excellent at generating high-resolution results. Diffusion models are very good at handling high-detail images, in contrast to other generative models like GANs, which may have trouble with high-resolution creation and display problems. This is crucial for fashion, as subtle details like fabric textures, zippers, and tiny accessories must be represented with visual clarity.

It is advantageous to train diffusion models on a range of data, including different clothing styles, materials, and accessories. Because of this exposure, the models are able to create stunning outcomes that accurately depict the diverse spectrum of fashion styles. Whether it is creating a durable leather jacket or a gorgeous evening gown, the model can ensure that every output is visually appealing and true to form.

3.2. Impact on sustainable fashion

One of the main reasons for waste in the fashion business is overproduction. Often, clothing companies produce more than they can sell, leaving unsold inventory that may end up in landfills. DiFashion may mitigate this by offering data-driven, demand-driven suggestions. By employing AI-generated recommendations to understand customer preferences, retailers can better predict what styles and quantities are in demand, removing the need for overproduction.

4. Systematic analysis of different outfit recommendation systems

Table 1: Systematic Analysis of Different Outfit Recommendation Systems

Sr no	Paper Title	Year	Dataset	Methodology	Findings	Limitations
1.	Effects of 3D Virtual “Try-On” on Online Sales and Customers’ Purchasing Experiences	2022	Fabric net	Data Collection: The study used actual sales data and customer interview data	Sales Increase: The average sales per customer increased by 14,000 won (approximately \$13 USD) after using the 3D virtual try-on.	Tactile Simulation: The VTO was limited by its inability to simulate the tactile sense or texture of fabrics accurately.
2.	FabricNet: Fiber Recognition Architecture Using Ensemble ConvNets	2021	Fabric net	Data Collection: Utilized the FabricNet dataset, including fabric images	High Accuracy: The model demonstrated significant accuracy in recognizing textile fibers across different conditions	Dataset Size: Limited size of the FabricNet dataset may restrict generalizability
3.	Adaptable Recommendation System for Outfit Selection with Deep Learning Approach	2024	Fabric net	Database Generator: Encodes the visual characteristics of garments.	The system effectively adapts to user preferences by leveraging both short-term and long-term memory.	The system’s initial performance depends on the accuracy of professional labeling to avoid random starts.
4.	Associate Design of Fashion Sketch and Pattern	2023	Fabric net	Proposed an associate design method to integrate fashion.	Generated fashion sketches accurately represented jean characteristics.	Some designers (10%) were still dissatisfied with generated sketches
5.	Analysis of Recommender System Using Generative Artificial Intelligence	2021	Fabric net	Feature Extraction: Employs a convolutional neural network (CNN).	The interactive adaptation process significantly narrows the gap between generic and user-specific clothing feature spaces.	Focuses solely on upper body clothing, limiting its applicability for full outfit aggregation.
6.	Multiple-Clothing Detection and Fashion Landmark Estimation Using a Single-Stage Detector	2021	Deep Fashion2	A fast, single-stage EfficientDet-based model for clothing detection and landmark estimation.	The method achieved 0.686 mAP in bounding box detection and 0.450 mAP in landmark estimation with an inference time of 42 ms	The model struggles with detecting landmarks when clothes are severely occluded or too small
7.	Text-Conditioned Outfit Recommendation with Hybrid Attention Layer	2024	Polyvore	A hybrid attention model using FashionCLIP to recommend outfits based on text descriptions.	Achieved state-of-the-art performance on outfit compatibility and retrieval tasks.	Struggles with conflicting text descriptions and partial outfit items
8.	Differentiated Fashion Recommendation Using Knowledge Graph and Data Augmentation	2019	Amazon Fashion	Data augmentation and knowledge graph to enhance recommendations for both active and inactive users.	Improved recommendation accuracy and alleviated cold start problem.	Increased complexity and time consumption in deeper search depths.
9.	Promoting Green Fashion Consumption Through Digital Nudges in Recommender Systems	2024	Custom dataset	Incremental clustering with user feedback and feature adaptation.	Nudges promoting sustainability and second-hand items increased green consumption, though ethical standards were less influential	The study's focus was limited to air and water conservation, lacking broader sustainability indicators.
10.	Aggregating Everyday Outfits by Incremental Clustering with Interactive User Adaptation	2021	Custom dataset	Incremental clustering with feedback and Siamese network adaptation.	Reduced user feedback and improved clustering accuracy through user-specific adaptation.	Dataset is limited to a fixed set of clothing over a year, and doesn't consider changes in wardrobe such as new purchases or disposals

With growing interest in sophisticated generative models like diffusion models, the field of fashion recommendation systems has seen significant change. Unlike conventional techniques like matrix factorization (MF), collaborative filtering (CF), and content-based filtering (CBF), these models provide individualized, high-fidelity, and original content, adding a new dimension to

recommendation systems. In this overview of the literature, we examine the main ideas and strategies of both conventional and contemporary techniques, concentrating on the relative benefits and drawbacks of diffusion models in relation to fashion advice.

5. Parameter analysis

Table 2: Comparative Analysis of Various Parameters of Outfit Recommendation System

Paper Title	Integration	Operational Time	Operational Cost	Efficiency	Data Security	Error Handling	Accuracy
[4]	AI	Moderate	Moderate	High	Moderate	Moderate	High
[5]	AI	Moderate	Low	High	Moderate	Moderate	Low
[6]	AI	Moderate	Moderate	High	Moderate	Moderate	Low
[7]	Human	Low	Low	High	Moderate	Moderate	High
[9]	Human	Moderate	Moderate	High	Moderate	Moderate	Low
[10]	Human	Moderate	Moderate	High	Moderate	Moderate	High

The main advantages of diffusion models over conventional techniques in fashion recommendation systems are shown by the parameter analysis. Diffusion models offer better novelty, personalization, and high-fidelity image production by regulating variables such as latent space dimensions, noise schedules, and personalization guidance. Traditional approaches, on the other hand, are computationally efficient but do not provide the same degree of creativity and user participation throughout processing.

6. Results

Fashion recommendation systems, like the DiFashion model, have effectively integrated diffusion models with noteworthy results in terms of user satisfaction, visual fidelity, and personalization. Compared to more traditional methods like collaborative filtering and content-based filtering, DiFashion was able to achieve a 30–40% increase in personalization accuracy by creating highly customized outfits that meet particular customer preferences like style and body type. Furthermore, the model produced realistic, high-resolution images of recommended clothing, significantly enhancing the visual outputs' quality as judged by both objective standards such as the structural similarity index (SSIM) and subjective user feedback. In overall, users were more satisfied and engaged with the recommendations, highlighting the value of diffusion models in creating distinctive and personalized fashion experiences.

7. Discussions

The use of diffusion models in fashion recommendation systems, more especially, the DiFashion model—represents a significant advancement over traditional techniques by addressing significant limitations related to customization, novelty, and visual quality. Unlike collaborative filtering and content-based filtering, which often generate generic and repetitive recommendations, diffusion models use high-dimensional latent spaces and sophisticated sampling techniques to generate unique and personalized outfits that appeal to individual user preferences. Additionally, by allowing users to more clearly perceive how different items of clothing combine to form full outfits, diffusion models' ability to produce high-fidelity visual outputs enhances user engagement. This generative capacity opens the door to more interesting and fulfilling design possibilities by encouraging experimentation with a range of solutions and adhering to evolving fashion industry trends.

8. Conclusion

In conclusion, the employment of diffusion models in fashion recommendation systems, such the DiFashion model, greatly enhances user satisfaction, visual fidelity, and customization when compared to traditional methods. Diffusion methods, which provide unique, high-quality clothing suggestions based on user preferences, overcome the primary shortcomings of collaborative filtering and content-based filtering. Ultimately, by promoting more user engagement and permitting continuous adaptation based on user feedback, this innovative approach enhances the dynamic and satisfying fashion shopping experience. Diffusion models are a possible new avenue for providing high-fidelity, customized, and imaginative recommendations as the fashion industry adopts technology more and more.

9. Future scope

Diffusion models in fashion recommendation systems have a bright future thanks to chances for greater customization through thorough user data integration and real-time trend monitoring. Virtual try-on experiences could be completely transformed by the possible integration with augmented reality (AR), which would make shopping more immersive and interesting. Extending applications to cross-domain suggestions may also result in comprehensive aesthetic solutions, and emphasizing sustainability may persuade clients to choose environmentally friendly options. Increasing computing efficiency will enable real-time suggestions and scalability, while establishing robust user feedback loops will provide continuous model adaptation to evolving preferences. Collectively, these innovations have the power to fundamentally transform.

References

- [1] J. Huang, R. Feris, Q. Chen, and S. Yan, 'Cross-domain image retrieval with a dual attribute-aware ranking network,' in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), Dec. 2015, pp. 1062–1070. <https://doi.org/10.1109/ICCV.2015.127>.
- [2] M. H. Kiapour, X. Han, S. Lazebnik, A. C. Berg, and T. L. Berg, 'Where to buy it: Matching street clothing photos in online shops,' in Proc. IEEE Int. Conf. Comput. Vis. (ICCV), Dec. 2015, pp. 3343–3351. <https://doi.org/10.1109/ICCV.2015.382>.
- [3] H. Chen, A. Gallagher, and B. Girod, 'Describing clothing by semantic attributes,' in Proc. Eur. Conf. Comput. Vis., Berlin, Germany: Springer, 2012, pp. 609–623. https://doi.org/10.1007/978-3-642-33712-3_44.

- [4] L. Bossard, "Apparel classification with style," in Proc. Asian Conf. Comput. Vis., Berlin, Germany: Springer, 2012, pp. 321–335. https://doi.org/10.1007/978-3-642-37447-0_25.
- [5] W.-H. Cheng, S. Song, C.-Y. Chen, S. Chusnul Hidayati, and J. Liu, Fashion meets computer vision: A survey,' 2020, arXiv:2003.13988. [Online]. Available: <http://arxiv.org/abs/2003.13988>
- [6] 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. Comput. Vis. Pattern Recognit. (CVPR), Jun. 2016, pp. 1096–1104. <https://doi.org/10.1109/CVPR.2016.124>.
- [7] Z. Liu, Fashion landmark detection in the wild,' in Proc. Eur. Conf. Comput. Vis., Cham, Switzerland: Springer, 2016, pp. 229–245. https://doi.org/10.1007/978-3-319-46475-6_15.
- [8] S. Yan, Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang, "Unconstrained fashion landmark detection via hierarchical recurrent transformer networks," in Proc. 25th ACM Int. Conf. Multimedia, Oct. 2017, pp. 172–180. <https://doi.org/10.1145/3123266.3123276>.
- [9] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, and A. W. M. Smeulders, "Selective search for object recognition," Int. J. Comput. Vis., vol. 104, no. 2, pp. 154–171, Sep. 2013. <https://doi.org/10.1007/s11263-013-0620-5>.
- [10] R. Girshick, J. Donahue, T. Darrell, and J. Malik, Rich feature hierarchies for accurate object detection and semantic segmentation,' in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., Jun. 2014, pp. 580–587. <https://doi.org/10.1109/CVPR.2014.81>.
- [11] L. Jiao, F. Zhang, F. Liu, S. Yang, L. Li, Z. Feng, and R. Qu, "A survey of deep learning-based object detection," IEEE Access, vol. 7, pp. 128837–128868, 2019. <https://doi.org/10.1109/ACCESS.2019.2939201>.
- [12] S. Ren, Faster R-CNN: Towards real-time object detection with region proposal networks,' in Proc. Adv. Neural Inf. Process. Syst., 2015, pp. 1–9.
- [13] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Jun. 2016, pp. 779–788. <https://doi.org/10.1109/CVPR.2016.91>.
- [14] M. Tan, R. Pang, and Q. V. Le, EfficientDet: Scalable and efficient object detection,' in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), Jun. 2020, pp. 10781–10790. <https://doi.org/10.1109/CVPR42600.2020.01079>.
- [15] S. Lee, S. Oh, C. Jung, and C. Kim, A global-local embedding module for fashion landmark detection,' in Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshop (ICCVW), Oct. 2019, pp. 3153–3156 <https://doi.org/10.1109/ICCVW.2019.00387>.