

# Fashionability-Enhancing Outfit Image Editing with Conditional Diffusion Models

Qice Qin  
Waseda University  
Tokyo, Japan

qinqice@toki.waseda.jp

Yuki Hirakawa  
ZOZO Research  
Tokyo, Japan

yuki.hirakawa@zozo.com

Ryotaro Shimizu  
ZOZO Research  
Tokyo, Japan

ryotaro.shimizu@zozo.com

Takuya Furusawa  
ZOZO Research  
Tokyo, Japan

takuya.furusawa@zozo.com

Edgar Simo-Serra  
Waseda University  
Tokyo, Japan

ess@waseda.jp

## Abstract

Image generation in the fashion domain has predominantly focused on preserving body characteristics or following input prompts, but little attention has been paid to improving the inherent fashionability of the output images. This paper presents a novel diffusion model-based approach that generates fashion images with improved fashionability while maintaining control over key attributes. Key components of our method include: 1) fashionability enhancement, which ensures that the generated images are more fashionable than the input; 2) preservation of body characteristics, encouraging the generated images to maintain the original shape and proportions of the input; and 3) automatic fashion optimization, which does not rely on manual input or external prompts. We also employ two methods to collect training data for guidance while generating and evaluating the images. In particular, we rate outfit images using fashionability scores annotated by multiple fashion experts through OpenSkill-based and five critical aspect-based pairwise comparisons. These methods provide complementary perspectives for assessing and improving the fashionability of the generated images. The experimental results show that our approach outperforms the baseline Fashion++ in generating images with superior fashionability, demonstrating its effectiveness in producing more stylish and appealing fashion images.

## 1. Introduction

The fashion industry is continuously evolving, with remarkable advancements in AI technology [4, 10, 18, 19, 37, 38, 40–43, 45, 47]. One of the challenges in fashion



Figure 1. Illustrations of minimal edits to enhance fashionability. 1) Modifying white pants to brown and adding floral patterns to the blouse (left). 2) Adding a belt to the outfit as an accessory (center). 3) Adjusting the shape and size of a loose-fitting red jumpsuit to a more form-fitting and tailored version (right).

e-commerce is the lack of expert advice that consumers would typically receive in physical stores, making it essential to develop features that support users’ decision-making. In particular, recent developments in virtual try-on systems [6–9, 20, 27, 34, 35, 39, 46, 48–50] and clothed human generation models [5, 12, 13, 21, 30], enable users to preview clothing on themselves without visiting physical stores. These models mainly focus on preserving body characteristics or following predefined prompts to generate fashion images [2, 28]. However, they often overlook enhancing the inherent fashionability of the generated images, which limits their effectiveness in supporting users unfamiliar with fashion.

To address this issue, we introduce an approach to enhance the fashionability of input images, helping users and designers make better fashion decisions. Figure 1 shows the concept of enhancing fashionability through minimal edits.<sup>1</sup> It showcases three pairs of images that demonstrate

<sup>1</sup>The base images (right) are sourced from the VITON-HD dataset [6].

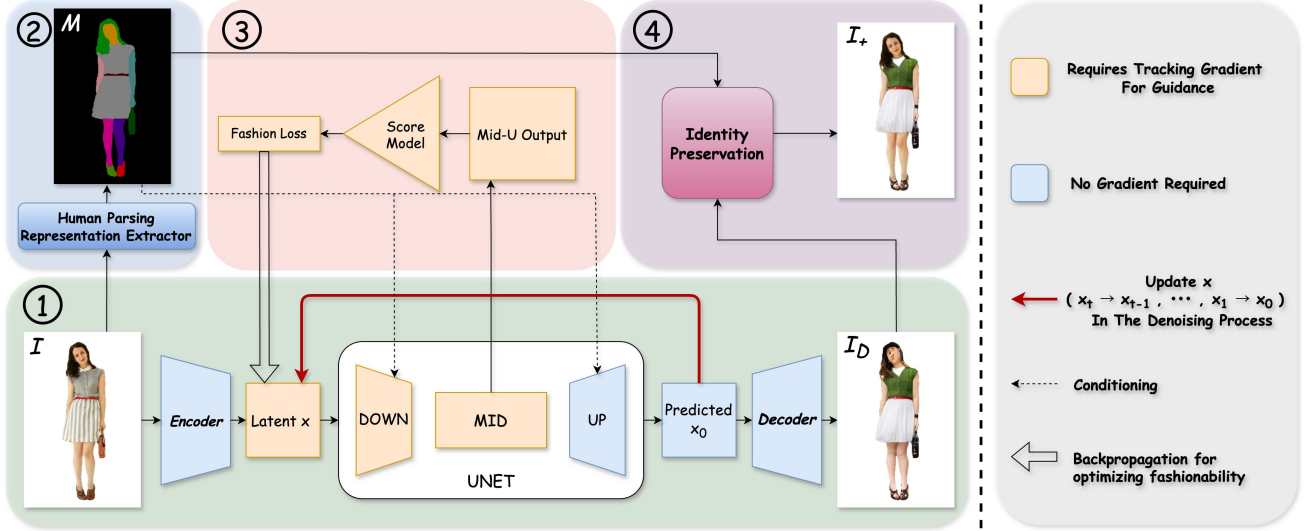


Figure 2. Overview of our proposed approach for conditional fashion image generation. Our approach consists of four main components: 1) a diffusion-based generation module that applies a diffusion process to iteratively refine the input image into a fashion-enhanced output while maintaining visual coherence; 2) a human parsing representation extractor that generates segmentation maps from the input image; 3) a mid-U classifier that processes mid-level UNet outputs to compute the fashion loss, which is fed back to the latent representation for fashionability optimization; and 4) an identity preservation process that combines the generated output with the segmentation map to restore the original subject’s head onto the new image.

how minor adjustments, such as modifying colors and patterns, adding accessories and changing clothing styles, can improve overall fashionability.

In this paper, we propose a novel approach to enhance the fashionability of generated images by leveraging intermediate latent features and structured conditioning during the denoising process of a latent diffusion model [36]. Our method incorporates a two-fold strategy: (1) utilizing an intermediate feature-based classifier to guide the diffusion process and (2) integrating segmentation map-based conditioning to ensure structural fidelity in the generated images.

First, We train a classifier on an expert-annotated dataset to predict fashionability scores and incorporate it into the generation process to produce images with enhanced fashionability. By introducing a guidance model trained on expert-annotated fashionability data, our approach enables more reliable generation compared to methods like Fashion++, which relies on artificially generated negative examples.

Additionally, we train a diffusion model conditioned on segmentation maps using images from WEAR [53]<sup>2</sup> to preserve the original clothing and body shape characteristics. Unlike previous methods [2, 28] that use detailed prompts describing garment styles and colors, our approach does not require specific prompts during training and inference. This allows our model to generate more fashionable images

without external user inputs. Finally, we use the diffusion pipeline augmented with Classifier Guidance. We calculate the fashionability loss of the latent features at each step using the trained classifier. This fashionability loss guides the diffusion model toward generating images with higher fashionability.

To assess the effectiveness of our method, we incorporate two additional evaluation models. The first is an OpenSkill [23]-based model, which leverages annotated pairwise comparisons to assign continuous fashionability scores, providing a quantitative evaluation framework. The second is a classifier that evaluates fashionability across five critical dimensions—*cleanliness*, *harmony*, *silhouette*, *styling*, and *trendiness*—offering a more comprehensive perspective on fashion perception.

Extensive experiments demonstrate that our method outperforms the baseline Fashion++ [17] in generating fashionable images. By combining a reliable, expert-annotated dataset with recent advancements in machine learning techniques, our research provides a robust framework for automatically enhancing the fashionability of generated images.

The main contributions of this study are summarized as follows:

- Our approach combines ControlNet with Classifier Guidance, allowing the model to generate highly fashionable images while preserving the input image’s body and clothing shape characteristics.

and the negative examples are made by InstructPix2Pix [3].

<sup>2</sup>The largest fashion outfit-sharing application in Japan.

- We develop two complementary fashionability prediction models to assess the fashionability of fashion images. The first model is trained on a dataset with absolute fashion scores annotated by fashion experts. The second is trained on expert evaluations across five critical aspects of fashionability, providing a comprehensive framework for evaluation.
- Our method undergoes evaluation, including quantitative metrics and visual quality assessment, demonstrating its superiority in generating highly fashionable images compared to the existing method.

## 2. Related Work

### 2.1. Fashion image generation

Fashion image generation has been a prominent area of research, and various methods have been proposed to create realistic and stylish fashion images. Hsiao et al. [17] introduced Fashion++, a data-driven approach that uses a BicycleGAN [52] framework to enhance the fashionability of existing outfits with minimal edits. Specifically, Fashion++ makes subtle but effective changes, such as modifying colors, adjusting textures, or removing unnecessary accessories, to create visually appealing outfits. The model is trained on a dataset of outfit images, using a loss function that combines reconstruction loss and feature matching loss to ensure the edits are both realistic and fashion-forward. As the sole baseline in this field, Fashion++ demonstrates the potential of leveraging generative models to improve clothing aesthetics in a data-driven manner. However, it has notable limitations, such as generating low-resolution images and relying on augmented negative samples (i.e., slightly altered versions of the original images) for training. This reliance on artificially generated data may overlook the subjective and multifaceted nature of fashionability, failing to fully capture its nuanced aspects. In contrast, our approach leverages real annotations in order to modify images with increased fashionability with higher fidelity.

On the other hand, Lampe et al. [28] proposed DiCTI, a diffusion-based clothing designer guided by textual input, focusing on generating fashion images based on user-provided descriptions. This method showcases the potential of diffusion models in creating diverse and stylish clothing designs but relies on detailed textual prompts. On the contrary, our approach can utilize diffusion models to improve fashionability without needing additional text prompts.

### 2.2. Fashionability Prediction

Fashionability prediction has been challenging due to its subjective nature and the numerous factors influencing fashion perception. Prior research has developed datasets by engaging human annotators in comparing pairs of images featuring individuals dressed in diverse outfits and selecting

the one deemed more aesthetically pleasing [1, 32]. Subsequently, these datasets were utilized for training and assessing deep learning models capable of evaluating fashion aesthetics.

Simo-Serra et al. [44] made significant contributions in this area by developing a model for analyzing the perception of fashionability. Their model considers various factors, including outfit styles, individual user characteristics, the background of the photograph, and the resulting fashionability score. It provides detailed suggestions to users, such as recommending specific garment adjustments or background changes to enhance their perceived fashionability.

Hirakawa et al. [16] conducted an in-depth analysis of the correlation between human judgments and predictions made by GPT-4 [33]. They demonstrated that while GPT-4 shows some correlation with human perception, GPT-4 still struggles to achieve human-level performance. This highlights the inherent difficulty in predicting fashionability, even for GPT-4 in zero-shot.

Motivated by these findings, we adopt an alternative approach for fashionability prediction that emphasizes precision and human alignment. We prepare small but carefully annotated datasets and train evaluation models on them, aiming to better capture the nuanced and subjective aspects of fashionability. Our approach avoids the pitfalls of the general-purpose model and enables us to evaluate the fashionability of images in a more reliable and comprehensive manner.

## 3. Methodology

Figure 2 illustrates the four main modules of our proposed method for conditional fashion image generation: **1) Diffusion-based generation module:** The input image is processed through a diffusion model to produce a new output image to enhance its fashionability. **2) Human parsing representation extractor:** Segmentation maps are generated from the input image to preserve the subject’s structural and proportional details. **3) Fashionability feedback loop:** The mid-U output from the UNet is evaluated by a mid-U classifier, which calculates a fashion loss. This loss is then used to refine the latent representation in the diffusion process, guiding the model towards generating images with higher fashionability. **4) Identity preservation:** To ensure the output retains the subject’s identity, the generated image and the segmentation map are combined to restore the original subject’s head onto the new image.

Our system’s modular design allows it to simultaneously enhance fashionability, maintain body characteristics, and preserve the subject’s identity, making it both practical and interpretable.




Figure 3. Illustration of the OpenSkill-based fashionability scoring process. The first part shows images ranked along an axis, representing their average fashionability scores with associated normal distributions. The second part depicts the annotation process, where human evaluators provide pairwise scores comparisons to determine which image is more fashionable. The third part reflects the updated scores and distributions of the images, adjusted based on the annotations.

### 3.1. Dataset and Preprocessing

Acquiring annotations for fashionability is inherently challenging due to its highly subjective and nuanced nature. Our study utilizes two datasets: the OpenSkill-based and 5-scores-based datasets, each tailored for specific approaches discussed in the following subsections. Each image in our datasets was carefully evaluated by annotators across multiple aspects, making the annotation process both labor-intensive and costly. As a result, building a large-scale dataset is impractical within this context. Despite the relatively small size of our datasets, their detailed and high-quality annotations ensure their utility in advancing the research on fashion image generation and fashionability evaluation. Figure 4 illustrates the construction and division of these datasets into training and testing subsets. This visualization provides a comprehensive overview of the datasets used throughout the study.

#### 3.1.1 OpenSkill-based Annotations

We adopt OpenSkill [23], a framework originally designed for estimating player skills in competitive games. OpenSkill models each player’s skill as a normal distribution characterized by a mean ( $\mu$ ) and standard deviation ( $\sigma$ ), which are updated based on the outcomes of pairwise matchups. In our case, we treat fashion snapshots as players and their aesthetic scores as skill ratings.

The annotation process consists of three main steps, as illustrated in Figure 3. Initially, the fashion snapshots are represented as normal distributions with prior scores. First, we conduct a matchup by sampling two snapshots based on the expected information gain from their comparison. Second, an annotator is randomly chosen to evaluate the pair and determine the more fashionable outfit. Third, the OpenSkill algorithm updates the parameters of both snapshots based on the annotator’s decision. These steps are iterated until the system converges to stable scores for all snapshots. Following a similar methodology to Kiapour et al. [24], which utilized TrueSkill [15], a predecessor of OpenSkill, we con-

structed image pairs for comparison by identifying the image with the lowest match rate and pairing it with a counterpart predicted to have a high likelihood of resulting in a draw, as determined by the OpenSkill scores. To ensure the accuracy of the annotations, we prepared two independent annotator groups and repeated the above procedure until Spearman’s rank correlation coefficient between the groups reached saturation. The final coefficient is 0.814, indicating that our dataset is highly reliable.<sup>3</sup>

Our dataset comprises 6,000 fashion snapshots from WEAR [53], annotated by hundreds of human evaluators. These snapshots underwent 219,488 pairwise comparisons, forming the foundation of our scoring process. Out of the 6,000 images, 5,000 were allocated to train the mid-U classifier to guide the generation process, and the remaining 1,000 images were reserved to train an OpenSkill-based model for fashionability prediction as well as to test the mid-U classifier. Additionally, 500 images from the original 5,000 training set were set aside to test the OpenSkill-based model.

#### 3.1.2 5-Dimension Annotations

Additionally, we utilized a second proprietary dataset of 831 images, each evaluated by annotators across five aspects: *cleanliness*, *harmony*, *silhouette*, *styling*, and *trendiness*, which were determined as essential by professionals with extensive fashion knowledge. The five dimensions and their evaluation criteria are as follows: **1) Cleanliness:** evaluates whether the outfit appears clean and well-maintained. **2) Harmony:** assesses whether the wearer’s hairstyle, makeup, and skin tone match the overall outfit. **3) Silhouette:** examines the balance of proportions and fit relative to the wearer’s body. **4) Styling:** measures the sense of styling in how the outfit is worn, including combining colors and patterns. **5) Trendiness:** determines whether the outfit reflects current fashion trends.

We collected annotations from experts using pairwise

<sup>3</sup>The details of our method for creating datasets are shown in [16].






Figure 4. Construction of the dataset. It shows how we divided them for each model’s training and testing.

comparisons, in which annotators were presented with pairs of images and asked to determine which image performed better in a specific aspect. This process resulted in a total of 180,548 pairwise comparison transactions, distributed as follows: Cleanliness (32,628), Harmony (35,265), Silhouette (34,833), Styling (34,461) and Trendiness (43,361). These transactions provide a robust and detailed foundation for modeling the nuanced evaluation of fashion elements across these dimensions.

We normalized each score so that it takes an integer value ranging from one to five and enables consistent comparisons across all aspects. Furthermore, we calculated the average score across the five aspects to derive an overall measure of fashionability. We rounded it to the nearest integer, thus obtaining a single, unified fashionability score for each image. The overall metric and individual ones allow us to score both the general impression and specific strengths or weaknesses of the fashionability of outfits.

This dataset was divided to train our 5-scores-based model that assesses fashionability along these dimensions, with 89.9% of images used for training and the remaining 10.1% of images used for testing. During the training process, the dataset was initially divided into training, validation, and testing subsets, with 80% used for training, 10% for validation, and 10% for testing. However, to fully utilize the limited dataset, the validation subset was merged into the training set for the final model training for 30 epochs.

### 3.2. Base Architecture

Our approach leverages ControlNet as the backbone model to incorporate additional control into the image generation process, ensuring the generated images adhere to specific desired attributes. We use segmentation maps as the primary conditioning input for ControlNet. These segmen-

tation maps are generated from fashion images by an out-of-box human parsing representation extractor [29], which helps capture the overall structure of the original images, such as layout and pose. Specifically, ControlNet is conditioned on segmentation maps  $\mathbf{M}$  of the input image  $\mathbf{I}_{\text{img}}$ . During training, the forward process incrementally adds noise to the input image  $\mathbf{I}_{\text{img}}$ , resulting in a noisy image  $\mathbf{x}_t$ , where  $t$  indicates the time step, the number of times noise is added. The network  $\epsilon_\theta$  is trained to predict the noise added to the noisy image  $\mathbf{x}_t$  using the time step  $t$  and the segmentation maps  $\mathbf{M}$  as conditions. The objective function for training ControlNet is defined as follows:

$$\mathcal{L}_{\text{control}} = \mathbb{E}_{\mathbf{I}_{\text{img}}, t, \mathbf{M}, \epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{I}})} \left[ \|\epsilon - \epsilon_\theta(\mathbf{x}_t, t, \mathbf{M})\|_2^2 \right], \quad (1)$$

where  $\mathcal{L}_{\text{control}}$  represents the overall training objective for the diffusion model. This objective guides the fine-tuning process of diffusion models with ControlNet to ensure that the generated images retain the original characteristics of clothing and body shape. In our implementation, we fine-tuned all parameters of ControlNet, leveraging the full model capacity to adapt it to our dataset.

### 3.3. Classifier Guidance

Integrating Classifier Guidance into the diffusion model is a crucial aspect of our approach. The Classifier Guidance helps steer the diffusion process toward generating images that are more fashionable by utilizing a pre-trained classifier that provides feedback at each step of the image generation. Specifically, we use Mid-U Guidance [22] rather than the original Classifier Guidance technique [11]. The key difference between Mid-U Guidance and the original Classifier Guidance is that, while the original technique operates directly in pixel space and requires gradients to be backpropagated through both the classification model and the decoder of a Variational Autoencoder [25] (for latent diffusion models), Mid-U Guidance leverages intermediate features from the UNet’s encoder. This enables more efficient guidance by significantly reducing both computation and memory usage. We chose Mid-U Guidance because it is more computationally efficient, allowing for faster processing while still providing high-quality guidance in the fashionability enhancement process.

During the diffusion process, the UNet model produces intermediate latent features, specifically from the mid-block of the network, denoted as  $\epsilon_{\text{mid}}(\mathbf{x}_{t,i})$ . These features are extracted and used as input for a pre-trained fashion classifier  $f_\phi$ . The classifier categorizes these intermediate features into three distinct fashionability classes: high, medium, and low. The dataset 3.1.1 used for this classification consists of 6,000 images annotated with continuous scores. We normalized each score to an integer value ranging from one to three and it is evenly divided into these three classes, with

2,000 images per class. This setup allows us to explicitly incorporate the fashionability score into the guidance process by leveraging the inherent order of the class labels. Unlike standard classification tasks, where class labels are independent, our approach assigns greater importance to higher fashionability scores to align to enhance fashionability.

The fashionability loss  $\mathcal{L}_{\text{fashion}}$  is defined as:

$$\mathcal{L}_{\text{fashion}} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^3 \frac{\exp(f_{\phi}(\epsilon_{\text{mid}}(\mathbf{x}_{t,i}))_j)}{\sum_{k=1}^3 \exp(f_{\phi}(\epsilon_{\text{mid}}(\mathbf{x}_{t,i}))_k)} \cdot j, \quad (2)$$

where  $f_{\phi}$  represents the fashion classifier, which outputs probabilities for three fashionability classes, and  $j$  is the class label (1: low, 2: medium, 3: high). Here, the softmax function is explicitly defined to provide a probability distribution over the fashionability classes. The weighted summation of class probabilities ensures that the loss prioritizes higher fashionability scores, effectively serving as a soft ordinal regression mechanism.

The computed fashionability loss is then back-propagated through the network, where the gradient of the fashionability loss with respect to the latent representations  $\mathbf{x}_{t,i}$  is then computed as:

$$\nabla_{\mathbf{x}_{t,i}} \mathcal{L}_{\text{fashion}} = \lambda \frac{\partial \mathcal{L}_{\text{fashion}}}{\partial \mathbf{x}_{t,i}}, \quad (3)$$

where  $\lambda$  is a scalar value representing the guidance loss scale. Finally, the latents are updated iteratively using the computed gradient and scaled by  $\sigma^2$ , where  $\sigma$  is the noise standard deviation at the current step  $i$ :

$$\mathbf{x}_{t,i+1} = \mathbf{x}_{t,i} - \sigma^2 \nabla_{\mathbf{x}_{t,i}} \mathcal{L}_{\text{fashion}}. \quad (4)$$

This iterative update ensures that the generated image aligns with the desired fashionability characteristics, as determined by the aesthetic loss function.

## 4. Experimental Results

### 4.1. Implementation Details

We used a pre-trained Stable Diffusion model v1-5 [36] to train ControlNet as a base model. The dataset used for training consisted of 14,983 images from Wear.jp, along with corresponding segmentation maps that annotate the fashion items present in each image, such as tops, bottoms, and accessories. The image resolution was set to  $512 \times 512$ , and the learning rate was fixed to  $10^{-6}$ . In equations (3) and (4),  $\lambda$ , the scaling factor for the guidance loss, was empirically set to 0.1, while  $\sigma$ , the noise standard deviation at each iteration, we used the UniPCMultistepScheduler [51] to determine the noise schedule. To optimize memory usage, we set the batch size and gradient accumulation steps

to one and four. Additionally, the training employed an 8-bit Adam optimizer and memory-efficient attention using xformers [31]. The model was trained for 93,000 steps on four Tesla V100-SXM2-16GB GPUs. For the Mid-U Classifier, we took as input features of dimension  $1280 \times 8 \times 8$ , which are intermediate outputs from the UNet model. The fashion classifier  $f_{\phi}$  mentioned in section 3.3 consists of a sequential network that begins with a  $3 \times 3$  convolutional layer, which takes an input of 1,280 channels and reduces it to 256 channels with padding of one, followed by a ReLU activation. A  $2 \times 2$  max-pooling layer succeeds this, and then a second  $3 \times 3$  convolutional layer that reduces 256 channels to 128 channels, also with padding of one, followed by another ReLU activation. The output is passed through an adaptive average pooling layer, which reduces the spatial dimensions to  $2 \times 2$ , and then flattened to a 512-dimensional vector. A fully connected layer then reduces this to 64 dimensions with a ReLU activation. Finally, an output layer maps the 64-dimensional vector to three outputs representing the three fashionability classes. The Mid-U Classifier was trained on the OpenSkill-based dataset with the train-test split shown in Figure 4, using a cross-entropy loss function. We trained the model for 50 epochs with a batch size of 32 and used the Adam optimizer with its default parameters, setting the learning rate to  $10^{-4}$ .

### 4.2. Fashionability Prediction and Evaluation Metrics

The fashionability prediction model provides quantitative evaluations of generated images to compare different fashion image generation methods. We developed two distinct models for fashionability prediction: **1) OpenSkill-based model**, trained on a dataset of 1,000 images categorized into high, medium, and low fashionability with a ResNet-50 [14] classifier pre-trained on ImageNet [26]. We emphasize that those images do not overlap with those used to train the classifier in the generative model. **2) 5-scores-based model**, trained on the other proprietary and confidential dataset of 831 images, each evaluated by experts across five aspects: trend, silhouette, sense, cleanliness, and dressing. A ResNet classifier was trained to predict the average scores as fashionability (very high, high, medium, low, and very low fashionability) across the five categories. While the OpenSkill-based model offers a broad evaluation based on a large dataset, the 5-scores-based model provides a detailed, expert-driven evaluation, ensuring a solid assessment of fashionability enhancement across diverse criteria. To evaluate the performance of our method, we employed several key metrics that reflect the quality and fashionability of the generated images: **1) Accuracy of Fashionability Prediction Models:** We measured the accuracy of our ResNet classifier trained on two datasets: one with 1,000 images



Figure 5. Qualitative comparison of our proposed approach and Fashion++ for the same input.

Evaluator	Method	Increased ( $\uparrow$ )	Decreased ( $\downarrow$ )
OpenSkill-based	Fashion++	22%	28%
	Ours	<b>56%</b>	<b>12%</b>
5-scores-based	Fashion++	10%	20%
	Ours	<b>26%</b>	<b>12%</b>

Table 1. Comparison of fashionability predictions results by OpenSkill-based model and 5-scores-based model on 50 images generated by Fashion++ and our approach with the same input images. The numbers indicate the ratio of images rated as having increased or decreased fashionability by each model.

and another with 831 images. These models provide a baseline for assessing the effectiveness of our model in predicting fashion scores. **2) Comparison of Fashionability Scores:** We compared the fashionability scores of images generated by our model and those generated by Fashion++. This comparison assessed the increase or decrease in fashionability scores assigned by two different prediction models. **3) Visual Quality of Generated Images:** We qualitatively assessed the visual quality of images generated by our model versus those generated by Fashion++. This metric focuses on the generated images’ clarity, detail, and overall aesthetic appeal. **4) User Study:** We conducted a straightforward user study to ask participants if the generated images by Fashion++ and our method are more fashionable than the original images. This study validates whether actual users perceive the generated images as more fashionable.

### 4.3. Quantitative Results

The ResNet classifier trained on 1,000 images achieved an accuracy of 60% in predicting high, medium, and low fashionability categories on a separate test set from the 5,000 images used to train the Mid-U classifier, outperforming the random chance accuracy of 33%. The model trained on 831 images with detailed fashion scores achieved an accuracy of 65%, which is notably higher than the random chance accuracy of 20%. We consider these accuracies reasonable, given the highly subjective nature of fashion and its inherent challenges even for human judgment. In addition, we conducted an evaluation of fashionability predictions for 50 images generated by two methods: Fashion++ and our proposed method. The results, as summarized in Table 1, indicate a significant improvement in the fashionability of images generated by our method. These results collectively demonstrate the model’s ability to effectively enhance the original images’ fashionability, surpassing the baseline method, Fashion++.

### 4.4. Qualitative Results

From Figure 5, it is evident that for images not seen during training, the results produced by Fashion++ tend to make the subjects appear thinner and often fail to generate the complete body shape accurately, with some instances even missing parts of the legs and feet. Additionally, the fabric textures and overall image quality of the clothing generated by Fashion++ are not very clear. In contrast, our method successfully maintains the original body characteristics while generating clear and detailed images of the clothing. This improvement can be attributed to using a supervised classifier trained on our annotated data, which

Participants	Method	Increased (↑)	Decreased (↓)
1	Fashion++	7	23
	Ours	<b>25</b>	<b>9</b>
2	Fashion++	11	30
	Ours	<b>31</b>	<b>12</b>
3	Fashion++	3	36
	Ours	<b>28</b>	<b>8</b>
4	Fashion++	9	26
	Ours	<b>20</b>	<b>6</b>
5	Fashion++	14	28
	Ours	<b>17</b>	<b>16</b>
Total	Fashion++	44	143
	Ours	<b>121</b>	<b>51</b>

Table 2. The results of our user study. The numbers indicate the number of images rated as having increased or decreased fashionability. 50 pairs between input and Fashion++’s output and 50 pairs between input and our method’s output are assigned to each participant. There are a total of 250 responses for each of these two kinds of comparison, respectively.

provides precise guidance during generation, and the high-quality image synthesis capability of Stable Diffusion 1.5. These factors enable our model to make subtle adjustments to enhance fashionability, such as changing the color of a belt, the style of clothing, and the overall outfit coordination. These improvements demonstrate the effectiveness of our model in producing high-quality and fashionable images.

#### 4.5. User Study

To evaluate and compare the perceived fashionability of images generated by Fashion++ and our proposed approach. We also conducted a simple user study, where we recruited five participants. Initially, they were presented with 50 sets of images, each comprising an original input image, an output image generated by Fashion++, and an output image generated by our method. We inquired from them about the number of images within these 50 sets where the Fashion++ output exhibited an increase in fashionability compared to the input images, as well as the number of images where our method’s output demonstrated an enhanced fashionability compared to the input images. In total, we collected 250 responses, and Table 2 summarizes the results of this study.

#### 4.6. Negative Results

As shown in Figure 6, our approach occasionally fails to generate consistent and anatomically accurate results. In example (a), the woman is holding a bag in front of her dress, with the strap encircling part of the garment, but the enclosed area displays a markedly different texture and color. In example (b), the man’s glasses, hanging on his chest, are rendered as a vague patch rather than realistically




Figure 6. Examples of failure cases of our approach. For each pair of images, the left image represents the input image, and the right one shows the generated image of our approach. (a) and (b) high-light inconsistencies within the generated garments, where certain areas exhibit significantly different patterns or textures compared to the rest of the outfit. (c) and (d) demonstrate issues in generating body parts, such as missing hands or incorrect orientation of legs, revealing limitations in preserving accurate anatomical details.

integrated. These issues suggest that the model struggles to process accessory-garment interactions and small accessories due to a lack of specific annotations and limited training data for such cases. In example (c), the woman’s hand merges with her pants, creating an unnatural white patch, while in example (d), the crossed legs are distorted, resulting in an implausible pose. These errors highlight the model’s challenges in maintaining realistic body structures for complex poses or overlapping body parts, likely due to ambiguous segmentation maps or insufficient pose diversity during training.

### 5. Limitations and Discussions

In this paper, we introduced a novel approach to enhance the fashionability of garment images by combining Classifier Guidance and ControlNet. Our method leverages a custom-trained classifier and a specialized ControlNet pipeline to generate fashionably improved images, addressing the limitations of existing methods. By training two distinct classifiers for fashionability prediction—one based on OpenSkill and another utilizing detailed expert ratings across multiple fashion aspects—we have established a robust framework for evaluating fashion images. Extensive evaluations demonstrate that our method consistently surpasses the existing approach in generating visually ap-



peeling and fashion-forward images.

While our model significantly improves fashionability and generates high-quality images, several limitations remain. First, fashionability is inherently subjective, and despite using expert-annotated data, the generated outputs may not fully align with individual user preferences. Second, though effective, the proprietary dataset used for training is limited in size and diversity, which could impact the model’s robustness in representing a wider range of styles and trends. Third, shading preservation on generated garments remains inconsistent, likely due to the model’s focus on enhancing fashionability rather than maintaining precise lighting consistency. Lastly, our method relies on segmentation maps that constrain garment dimensions to the original body proportions, limiting flexibility in modifying shapes and designs. Addressing these limitations in future work could further enhance the versatility and applicability of our approach in the fashion industry, contributing to its evolution in both virtual and real-world contexts.

## Acknowledgments

We thank Takuma Nakamura, Kazuya Morishita, and Sai HtaungKham for their support for data collection.

## References

- [1] Eli Alshan, Sharon Alpert, Assaf Neuberger, Nathaniel Buis, and Eduard Oks. Learning fashion by simulated human supervision. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2401–24016, 2018. [3](#)
- [2] Alberto Baldrati, Davide Morelli, Giuseppe Cartella, Marcella Cornia, Marco Bertini, and Rita Cucchiara. Multimodal garment designer: Human-centric latent diffusion models for fashion image editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23393–23402, 2023. [1](#), [2](#)
- [3] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. [2](#)
- [4] Qianqian Chen, Tianyi Zhang, Maowen Nie, Zheng Wang, Shihao Xu, Wei Shi, and Zhao Cao. Fashion-gpt: Integrating llms with fashion retrieval system. In *Proceedings of the 1st Workshop on Large Generative Models Meet Multimodal Applications*, pages 69–78, 2023. [1](#)
- [5] Weifeng Chen, Tao Gu, Yuhao Xu, and Arlene Chen. Magic clothing: Controllable garment-driven image synthesis. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 6939–6948, 2024. [1](#)
- [6] Seunghwan Choi, Sunghyun Park, Minsoo Lee, and Jaegul Choo. Viton-hd: High-resolution virtual try-on via misalignment-aware normalization. In *Proc. of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2021. [1](#)
- [7] Yisol Choi, Sangkyung Kwak, Kyungmin Lee, Hyungwon Choi, and Jinwoo Shin. Improving diffusion models for authentic virtual try-on in the wild. In *Computer Vision – ECCV 2024*, pages 206–235, 2025. [1](#)
- [8] Zheng Chong, Xiao Dong, Haoxiang Li, Shiyue Zhang, Wenqing Zhang, Xujie Zhang, Hanqing Zhao, and Xiaodan Liang. CatVTON: Concatenation is all you need for virtual try-on with diffusion models. *arXiv preprint arXiv:2407.15886*, 2024. [1](#)
- [9] Aiyu Cui, Jay Mahajan, Viraj Shah, Preeti Gomathinayagam, Chang Liu, and Svetlana Lazebnik. Street tryon: Learning in-the-wild virtual try-on from unpaired person images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 8235–8239, June 2024. [1](#)
- [10] Yashar Deldjoo, Fatemeh Nazary, Arnau Ramisa, Julian McAuley, Giovanni Pellegrini, Alejandro Bellogin, and Tommaso Di Noia. A review of modern fashion recommender systems. *ACM Comput. Surv.*, 56(4), 2023. [1](#)
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [5](#)
- [12] Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen Change Loy, Wayne Wu, and Ziwei Liu. Stylegan-human: A data-centric odyssey of human generation. In *European Conference on Computer Vision*, pages 1–19. Springer, 2022. [1](#)
- [13] Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Wayne Wu, and Ziwei Liu. Unitedhuman: Harnessing multi-source data for high-resolution human generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7301–7311, 2023. [1](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [6](#)
- [15] Ralf Herbrich, Tom Minka, and Thore Graepel. Trueskill™: a bayesian skill rating system. *Advances in neural information processing systems*, 19, 2006. [4](#)
- [16] Yuki Hirakawa, Takashi Wada, Kazuya Morishita, Ryotaro Shimizu, Takuya Furusawa, Sai Htaung Kham, and Yuki Saito. An empirical analysis of gpt-4v’s performance on fashion aesthetic evaluation. In *SIGGRAPH Asia 2024 Technical Communications*, SA ’24, New York, NY, USA, 2024. Association for Computing Machinery. [3](#), [4](#)
- [17] Wei-Lin Hsiao, Isay Katsman, Chao-Yuan Wu, Devi Parikh, and Kristen Grauman. Fashion++: Minimal edits for outfit improvement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5047–5056, 2019. [2](#), [3](#)
- [18] Naoto Inoue, Edgar Simo-Serra, Toshihiko Yamasaki, and Hiroshi Ishikawa. Multi-label fashion image classification with minimal human supervision. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*, pages 2261–2267, 2017. [1](#)
- [19] Sk Maidul Islam, Subhankar Joardar, and Arif Ahmed Sekh. A survey on fashion image retrieval. *ACM Computing Surveys*, 56(6), 2024. [1](#)

- [20] Tasin Islam, Alina Miron, Xiaohui Liu, and Yongmin Li. Deep learning in virtual try-on: A comprehensive survey. *IEEE Access*, 12:29475–29502, 2024. 1
- [21] Yuming Jiang, Shuai Yang, Haonan Qiu, Wayne Wu, Chen Change Loy, and Ziwei Liu. Text2human: Text-driven controllable human image generation. *ACM Transactions on Graphics (TOG)*, 41(4):1–11, 2022. 1
- [22] Jonathan Whitaker. Mid-U Guidance: Fast Classifier Guidance for Latent Diffusion Models. Retrieved from <https://wandb.ai/johnnowhitaker/midu-guidance/reports/Mid-U-Guidance-Fast-Classifier-Guidance-for-Latent-Diffusion-Models--VmlldzozMjg0NzA1?galleryTag=stable-diffusion>. Accessed December 1, 2024, 2024. 5
- [23] Vivek Joshy. Openskill: A faster asymmetric multi-team, multiplayer rating system. *Journal of Open Source Software*, 9(93):5901, 2024. 2, 4
- [24] M Hadi Kiapour, Kota Yamaguchi, Alexander C Berg, and Tamara L Berg. Hipster wars: Discovering elements of fashion styles. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*, pages 472–488. Springer, 2014. 4
- [25] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 5
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 6
- [27] Gaurav Kuppa, Andrew Jong, Xin Liu, Ziwei Liu, and Teng-Sheng Moh. Shineon: Illuminating design choices for practical video-based virtual clothing try-on. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops*, pages 191–200, 2021. 1
- [28] Ajda Lampe, Julija Stopar, Deepak K Jain, Shinichiro Omachi, Peter Peer, and Vitomir Štruc. Dicti: Diffusion-based clothing designer via text-guided input. In *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition*, pages 1–9. IEEE, 2024. 1, 2, 3
- [29] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Self-correction for human parsing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):3260–3271, 2020. 5
- [30] Shikai Li, Jianglin Fu, Kaiyuan Liu, Wentao Wang, Kwan-Yee Lin, and Wayne Wu. Cosmicman: A text-to-image foundation model for humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6955–6965, 2024. 1
- [31] Meta Platforms, Inc. Xformers. Retrieved from <https://facebookresearch.github.io/xformers/>. Accessed December 1, 2024, 2021. 6
- [32] Assaf Neuberger, Sharon Alpert, Eli Alshan, Nati Bubis, and Eduard Oks. Learning fashion traits with label uncertainty. In *CVPR 2018*, 2018. 3
- [33] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 3
- [34] Kaicheng Pang, Xingxing Zou, and Waikeng Wong. Learning visual body-shape-aware embeddings for fashion compatibility. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 8056–8065, 2024. 1
- [35] Hamza Rawal, Muhammad Junaid Ahmad, and Farooq Zaman. Gc-vton: Predicting globally consistent and occlusion aware local flows with neighborhood integrity preservation for virtual try-on. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 5251–5260, 2024. 1
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, June 2022. 2, 6
- [37] A. Rubio, LongLong Yu, E. Simo-Serra, and F. Moreno-Noguer. Multi-modal joint embedding for fashion product retrieval. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, pages 400–404, 2017. 1
- [38] Yuki Saito, Takuma Nakamura, Hirotaka Hachiya, and Kenji Fukumizu. Exchangeable deep neural networks for set-to-set matching and learning. In *Computer Vision – ECCV 2020*, pages 626–646. Springer International Publishing, 2020. 1
- [39] Chengji Shen, Zhenjiang Liu, Xin Gao, Zunlei Feng, and Mingli Song. Self-adaptive clothing mapping based virtual try-on. *ACM Transactions on Multimedia Computing Communications and Applications*, 20(3), 2023. 1
- [40] Ryotaro Shimizu, Megumi Matsutani, and Masayuki Goto. An explainable recommendation framework based on an improved knowledge graph attention network with massive volumes of side information. *Knowledge-Based Systems*, 239:107970, 2022. 1
- [41] Ryotaro Shimizu, Takuma Nakamura, and Masayuki Goto. Fashion-specific ambiguous expression interpretation with partial visual-semantic embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3497–3502, 2023. 1
- [42] Ryotaro Shimizu, Yuki Saito, Megumi Matsutani, and Masayuki Goto. Fashion intelligence system: An outfit interpretation utilizing images and rich abstract tags. *Expert Systems with Applications*, 213:119167, 2023. 1
- [43] Ryotaro Shimizu, Yu Wang, Masanari Kimura, Yuki Hirakawa, Takashi Wada, Yuki Saito, and Julian McAuley. A fashion item recommendation model in hyperbolic space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 8377–8383, 2024. 1
- [44] Edgar Simo-Serra, Sanja Fidler, Francesc Moreno-Noguer, and Raquel Urtasun. Neuroaesthetics in fashion: Modeling the perception of fashionability. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 869–877, 2015. 3
- [45] Yuxin Tian, Shawn Newsam, and Kofi Boakye. Fashion image retrieval with text feedback by additive attention compositional learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1011–1021, 2023. 1

- [46] Chiang Tseng, Chieh-Yun Chen, and Hong-Han Shuai. Artifact does matter! low-artifact high-resolution virtual try-on via diffusion-based warp-and-fuse consistent texture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 8240–8244, June 2024. 1
- [47] Yongquan Wan, Wenhai Wang, Guobing Zou, and Bofeng Zhang. Cross-modal feature alignment and fusion for composed image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 8384–8388, 2024. 1
- [48] Zhening Xing, Yuchen Wu, Si Liu, Shangzhe Di, and Huimin Ma. Virtual try-on with garment self-occlusion conditions. *IEEE Transactions on Multimedia*, 25:7323–7336, 2023. 1
- [49] Xu Yang, Changxing Ding, Zhibin Hong, Junhao Huang, Jin Tao, and Xiangmin Xu. Texture-preserving diffusion models for high-fidelity virtual try-on. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7017–7026, 2024. 1
- [50] Shufang Zhang, Minxue Ni, Shuai Chen, Lei Wang, Wenxin Ding, and Yuhong Liu. A two-stage personalized virtual try-on framework with shape control and texture guidance. *IEEE Transactions on Multimedia*, 26:10225–10236, 2024. 1
- [51] Wenliang Zhao, Lujia Bai, Yongming Rao, Jie Zhou, and Jiwen Lu. Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024. 6
- [52] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. *Advances in neural Information Processing Systems*, 30, 2017. 3
- [53] ZOZO, Inc. WEAR. Retrieved from <https://wear.jp/>. Accessed December 1, 2024, 2024. 2, 4