

Multimodal Garment Designer: Human-Centric Latent Diffusion Models for Fashion Image Editing

Alberto Baldrati^{1,3,*}, Davide Morelli^{2,3,*}, Giuseppe Cartella², Marcella Cornia²,
 Marco Bertini¹, Rita Cucchiara^{2,4}

¹University of Florence, Italy ²University of Modena and Reggio Emilia, Italy

³University of Pisa, Italy ⁴IIT-CNR, Italy

¹{name.surname}@unifi.it

²{name.surname}@unimore.it



Figure 1. In this work, we propose a novel multimodal garment designer framework based on latent diffusion models that can generate a novel fashion image conditioned on text, human keypoints, and a garment sketch.

Abstract

Fashion illustration is used by designers to communicate their vision and to bring the design idea from conceptualization to realization, showing how clothes interact with the human body. In this context, computer vision can thus be used to improve the fashion design process. Differently from previous works that mainly focused on the virtual try-on of garments, we propose the task of multimodal-conditioned fashion image editing, guiding the generation of human-centric fashion images by following multimodal prompts, such as text, human body poses, and garment sketches. We tackle this problem by proposing a new architecture based on latent diffusion models, an approach that has not been used before in the fashion domain. Given the lack of existing datasets suitable for the task, we also extend two existing fashion datasets, namely Dress Code and VITON-HD, with multimodal annotations collected in a semi-automatic manner. Experimental re-

sults on these new datasets demonstrate the effectiveness of our proposal, both in terms of realism and coherence with the given multimodal inputs. Source code and collected multimodal annotations are publicly available at: <https://github.com/aimagelab/multimodal-garment-designer>.

1. Introduction

Computer Vision research has always paid much attention both to the human person and to fashion-related problems, especially working on the recognition and retrieval of clothing items [11, 22], the recommendation of similar garments [8, 18, 39], and the virtual try-on of clothes and accessories [7, 13, 27, 28, 47, 52]. In the last years, some research efforts have been dedicated to the text-conditioned image editing task where, given a model image and a textual description of a garment, the goal is to generate the

*Equal contribution.

input model wearing a new clothing item corresponding to the given textual description. In this context, only a few works [19, 33, 56] have been proposed, exclusively employing GAN-based approaches for the generative step.

Recently, diffusion models [10, 17, 30, 41] have attracted more and more attention due to their outstanding generation capabilities, allowing the improvement of a variety of downstream tasks in several domains, while their applicability to the fashion domain is still unexplored. Many different solutions have been introduced and can roughly be identified based on the denoising conditions used to guide the diffusion process, which can enable greater control of the synthesized output. A particular type of diffusion model has been proposed in [37] that, instead of applying the diffusion process in the pixel space, defines the forward and the reverse processes in the latent space of a pre-trained autoencoder, becoming one of the leading choices thanks to its reduced computational cost. Although this solution can generate highly realistic images, it does not perform well in human-centric generation tasks and can not deal with multiple conditioning signals to guide the generation phase.

In this work, we address an extended and more general framework and define the new task of *multimodal-conditioned fashion image editing*, which allows guiding the generative process via multimodal prompts while preserving the identity and body shape of a given person (Fig. 1). To tackle this task, we introduce a new architecture, called Multimodal Garment Designer (MGD), that emulates the process of a designer conceiving a new garment on a model shape, based on preliminary indications provided through a textual sentence or a garment sketch. In particular, starting from Stable Diffusion [37], we propose a denoising network that can be conditioned by multiple modalities and also takes into account the pose consistency between input and generated images, thus improving the effectiveness of human-centric diffusion models.

To address the newly proposed task, we present a semi-automatic framework to extend existing datasets with multimodal data. Specifically, we start from two famous virtual try-on datasets (*i.e.* Dress Code [28] and VITON-HD [7]) and extend them with textual descriptions and garment sketches. Experimental results on the two proposed multimodal fashion benchmarks show both quantitatively and qualitatively that our proposed architecture generates high-quality images based on the given multimodal inputs and outperforms all considered competitors and baselines, also according to human evaluations.

To sum up, our contributions are as follows: (1) We propose a novel task of multimodal-conditioned fashion image editing, which entails the use of multimodal data to guide the generation. (2) We introduce a new human-centric generative architecture based on latent diffusion models, capable of following multimodal prompts while preserving the

model’s characteristics. (3) To tackle the new task, we extend two existing fashion datasets with textual sentences and garment sketches devising a semi-automatic annotation framework. (4) Extensive experiments demonstrate that the proposed approach outperforms other competitors in terms of realism and coherence with multimodal inputs.

2. Related Work

Text-Guided Image Generation. Creating an image that faithfully reflects the provided textual prompt is the goal of text-to-image synthesis. In this context, early approaches were based on GANs [45, 51, 53, 55], while most recent solutions exploit the effectiveness of diffusion models [31, 35, 37]. In the fashion domain, only a few attempts of text-to-image synthesis have been proposed [19, 33, 56]. Specifically, Zhu *et al.* [56] presented a GAN-based solution that generates the final image conditioned on both textual descriptions and semantic layouts. A different approach is the one introduced in [33], where a latent code regularization technique is employed to augment the GAN inversion process by exploiting CLIP textual embeddings [34] to guide the image editing process. Instead, Jiang *et al.* [19] proposed an architecture that synthesizes full-body images by mapping the textual descriptions of clothing items into one-hot vectors, limiting however the expressiveness capability of the conditioning signal.

Multimodal Image Generation with Diffusion Models. A related line of works aims to condition existing diffusion models on different modalities thus enabling greater control over the generation process [5, 6, 25, 29, 48]. For example, Choi *et al.* [6] proposed to refine the generative process of an unconditional denoising diffusion probabilistic model [30] by matching each latent variable with the given reference image. On a different line, the approach introduced in [25] adds noise to a stroke-based input and applies the reverse stochastic differential equation to synthesize images, without additional training. Wang *et al.* [48], instead, proposed to learn a highly semantic latent space and perform conditional finetuning for each downstream task to map the guidance signals to the pre-trained space. Other recent works proposed to add sketches as additional conditioning signals, either concatenating them with the model input [5] or training an MLP-based edge predictor to map latent features to spatial maps [46].

Among contemporary works that aim to condition pre-trained latent diffusion models, ControlNet [54] proposes to extend the Stable Diffusion model [37] with an additional conditioning input. This process involves creating two versions of the original model’s weights: one that remains fixed and unchanged (locked copy) and another that can be updated during training (trainable copy). The purpose of this is to allow the trainable version to learn the newly introduced condition while the locked version retains the origi-

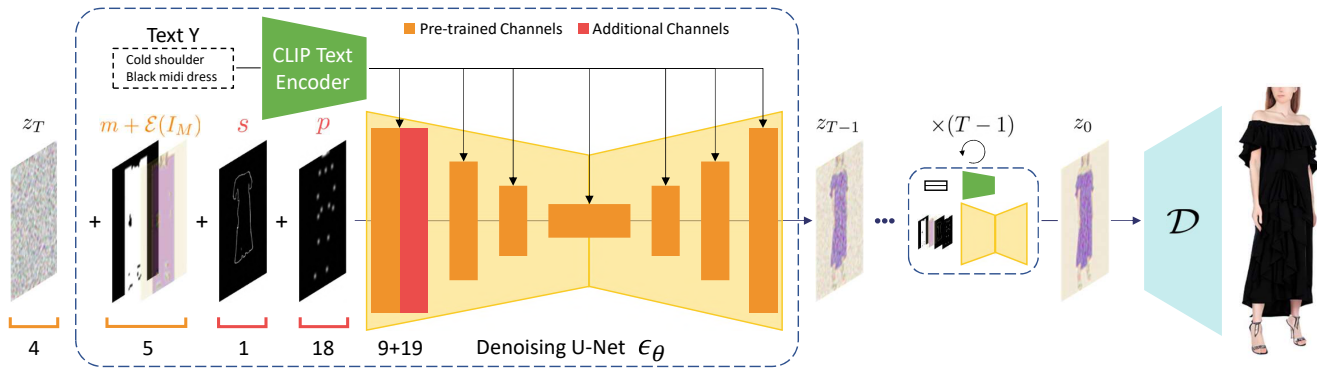


Figure 2. Overview of the proposed Multimodal Garment Designer (MGD), a human-centric latent diffusion model conditioned on multiple modalities (*i.e.* text, human pose, and garment sketch).

nal model knowledge. On the other hand, T2I-Adapter [29] learns modality-specific adapter modules that enable Stable Diffusion conditioning on new modalities.

In contrast, we focus on the fashion domain and propose a human-centric architecture based on latent diffusion models that directly exploits the conditioning of textual sentences and other modalities such as human body poses and garment sketches.

3. Proposed Method

In this section, we propose a novel task to automatically edit a human-centric fashion image conditioned on multiple modalities. Specifically, given the model image $I \in \mathbb{R}^{H \times W \times 3}$, its pose map $P \in \mathbb{R}^{H \times W \times 18}$ where the channels represent the human keypoints, a textual description Y of a garment, and a sketch of the same $S \in \mathbb{R}^{H \times W \times 1}$, we want to generate a new image $\tilde{I} \in \mathbb{R}^{H \times W \times 3}$ that retains the information of the input model while substituting the target garment according to the multimodal inputs. To tackle the task, we propose a novel latent diffusion approach, called Multimodal Garment Designer (MGD), that can effectively combine multimodal information when generating the new image \tilde{I} . Our proposed architecture is a general framework that can be easily extended to other modalities such as texture and 3d information. We strongly believe this task can foster research in the field and enhance the design process of new fashion items with greater customization. An overview of our model is shown in Fig. 2.

3.1. Preliminaries

While diffusion models [41] are latent variable architectures that work in the same dimensionality of the data (*i.e.* in the pixel space), latent diffusion models (LDMs) [37] operate in the latent space of a pre-trained autoencoder achieving higher computational efficiency while preserving the generation quality. In our work, we leverage the Stable Diffusion model [37], a text-to-image implementation of LDMs as a starting point to perform multimodal condition-

ing for human-centric fashion image editing. Stable Diffusion is composed of an autoencoder with an encoder \mathcal{E} and a decoder \mathcal{D} , a text-time-conditional U-Net denoising model ϵ_θ , and a CLIP-based text encoder T_E taking as input a text Y . The encoder \mathcal{E} compresses an image I into a lower-dimensional latent space defined in $\mathbb{R}^{h \times w \times 4}$, where $h = H/8$ and $w = W/8$. The decoder \mathcal{D} performs the opposite operation, decoding a latent variable into the pixel space. For the sake of clarity, we define the ϵ_θ convolutional input (*i.e.* z_t in this case) as spatial input γ because of the property of convolutions to preserve the spatial structure and the attention conditioning input as ψ . The denoising network ϵ_θ is trained according to the following loss:

$$L = \mathbb{E}_{\mathcal{E}(I), Y, \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_\theta(\gamma, \psi)\|_2^2], \quad (1)$$

where t is the diffusing time step, $\gamma = z_t$, $\psi = [t; T_E(Y)]$, and $\epsilon \sim \mathcal{N}(0, 1)$ is the Gaussian noise added to $\mathcal{E}(I)$.

3.2. Human-Centric Image Editing

Our task aims to generate a new image \tilde{I} , by replacing in the input image I the target garment using multimodal inputs, while preserving the model’s identity and physical characteristics. As a natural consequence, this task can be identified as a particular type of inpainting tailored for human body data. Instead of using a standard text-to-image model, we perform inpainting concatenating along the channel dimension of the denoising network input z_t an encoded masked image $\mathcal{E}(I_M)$ and the relative resized binary inpainting mask $m \in \{0, 1\}^{h \times w \times 1}$, which is derived from the original inpainting mask $M \in \{0, 1\}^{H \times W \times 1}$. Since here, the spatial input of the denoising network is $\gamma = [z_t; m; \mathcal{E}(I_M)]$, $\gamma \in \mathbb{R}^{h \times w \times 9}$. Thanks to the fully convolutional nature of the encoder \mathcal{E} and the decoder \mathcal{D} , this LDMs-based architecture can preserve the spatial information in the latent space. Exploiting this feature, our method can thus optionally add conditioning constraints to the generation. In particular, we propose to add two generation constraints in addition to the textual information: the

model pose map P to preserve the original human pose of the input model and the garment sketch S to allow the final users to better condition the garment generation process.

Pose Map Conditioning. In most cases [21, 24, 44], inpainting is performed with the objective of either removing or entirely replacing the content of the masked region. However, in our task, we aim to remove all information regarding the garment worn by the model while preserving the model’s body information and identity. Thus, we propose to improve the garment inpainting process by using the bounding box of the segmentation mask along with pose map information representing body keypoints. This approach enables the preservation of the model’s physical characteristics in the masked region while allowing the inpainting of garments with different shapes. Differently from conventional inpainting techniques, we focus on selectively retaining and discarding specific information within the masked region to achieve the desired outcome. To enhance the performance of the denoising network with human body keypoints, we modify the first convolution layer of the network by adding 18 additional channels, one for each keypoint. Adding new inputs usually would require retraining the model from scratch, thus consuming time, data, and resources, especially in the case of data-hungry models like the diffusion ones. Therefore, we propose to extend the kernels of the pre-trained input layer of the denoising network with randomly initialized weights sampled from a uniform distribution [14] and retrain the whole network. This consistently reduces the number of training steps and enables training with less data. Our experiments show that such improvement enhances the consistency of the body information between the generated image and the original one.

Incorporating Sketches. Fully describing a garment using only textual descriptions is a challenging task due to the complexity and ambiguity of natural language. While text can convey specific attributes like style, color, and patterns of a garment, it may not provide sufficient information about its spatial characteristics, such as shape and size. This limitation can hinder the customization of the generated clothing item other than the ability to accurately match the user’s intended style. Therefore, we propose to leverage garment sketches to enrich the textual input with additional spatial fine-grained details. We achieve this following the same approach described for pose map conditioning. The final spatial input of our denoising network is $\gamma = [z_t; m; \mathcal{E}(I_M); p; s]$, $[p; s] \in \mathbb{R}^{h \times w \times (18+1)}$, p and s are obtained by resizing P and S to match the latent space dimensions. In the case of sketches, we only condition the early steps of the denoising process as the final steps have little influence on the shapes [2].

Mask Composition. To preserve the model identity when performing human-centric inpainting, we perform mask composition as the final step of the proposed approach.

Defining $\hat{I} = \mathcal{D}(z_0) \in \mathbb{R}^{H \times W \times 3}$ as the output of the decoder \mathcal{D} and $M_{\text{head}} \in \{0, 1\}^{H \times W \times 1}$ as the model face binary mask of the image I , the final output image \tilde{I} is obtained as follows: $\tilde{I} = M_{\text{head}} \odot I + (1 - M_{\text{head}}) \odot \hat{I}$, where \odot denotes the element-wise multiplication operator.

3.3. Training and Inference

As in standard latent diffusion models, given an encoded input $z = \mathcal{E}(I)$, the proposed denoising network is trained to predict the noise stochastically added to z . The corresponding objective function can be specified as

$$L = \mathbb{E}_{\mathcal{E}(I), Y, \epsilon \sim \mathcal{N}(0,1), t, \mathcal{E}(I_M), m, p, s} [\|\epsilon - \epsilon_\theta(\gamma, \psi)\|_2^2], \quad (2)$$

where $\gamma = [z_t; m; \mathcal{E}(I_M); p; s]$ and $\psi = [t; T_E(Y)]$.

Classifier-Free Guidance. Classifier-free guidance is an inference technique that requires the denoising network to work both conditioned and unconditioned. This method modifies the unconditional model predicted noise moving it toward the conditioned one. Specifically, the predicted diffusion process at time t , given the generic condition c , is computed as follows:

$$\hat{\epsilon}_\theta(z_t|c) = \epsilon_\theta(z_t|\emptyset) + \alpha \cdot (\epsilon_\theta(z_t|c) - \epsilon_\theta(z_t|\emptyset)), \quad (3)$$

where $\epsilon_\theta(z_t|c)$ is the predicted noise at time t given the condition c , $\epsilon_\theta(z_t|\emptyset)$ is the predicted noise at time t given the null condition, and the guidance scale α controls the degree of extrapolation towards the condition.

Since our model deals with three conditions (*i.e.* text, pose map, and sketch), we use the fast variant multi-condition classifier-free guidance proposed in [1]. Instead of performing the classifier-free guidance according to each condition probability, it computes the direction of the joint probability of all the conditions $\Delta_{\text{joint}}^t = \epsilon_\theta(z_t|\{c_i\}_{i=1}^N) - \epsilon_\theta(z_t|\emptyset)$:

$$\hat{\epsilon}_\theta(z_t|\{c_i\}_{i=1}^N) = \epsilon_\theta(z_t|\emptyset) + \alpha \cdot \Delta_{\text{joint}}^t. \quad (4)$$

This reduces the number of feed-forward executions from $N + 1$ to 2.

Unconditional Training. Ensuring the ability of the denoising model to work both with and without conditions is achieved by replacing at training time the condition with a null one according to a fixed probability. This approach allows the model to learn from both conditional and unconditional samples, resulting in improved mode coverage and sample fidelity. Moreover, this technique also allows the model to optionally use the control signals at prediction time. Since our approach considers multiple conditions, we propose to extend the input masking to each condition independently. Experiments show that tuning this parameter can effectively affect the quality of the final result.

4. Collecting Multimodal Fashion Datasets

Currently available datasets for fashion image generation often contain low-resolution images and lack all the required multimodal information needed to perform the task previously described. For this reason, the collection of new multimodal datasets for the fashion domain plays a crucial role to advance research in the field. To this aim, we start from two recent high-resolution fashion datasets introduced for the virtual try-on task, namely Dress Code [28] and VITON-HD [7], and extend them with textual sentences and garment sketches. Both datasets include image pairs with a resolution of 1024×768 , each composed of a garment image and a reference model wearing the given fashion item. In this section, we introduce a framework to semi-automatically annotate fashion images with multimodal information and provide a complete description of how to enrich Dress Code and VITON-HD with garment-related text and sketches. We call our extended versions of these datasets Dress Code Multimodal and VITON-HD Multimodal, respectively. Sample images and multimodal data of the collected datasets can be found in Fig. 3.

4.1. Dataset Collection and Annotation

Data Preparation. We start the annotation from the Dress Code dataset, which contains more than 53k model-garment pairs of multiple categories. As a first step, we need to associate each garment with a textual description containing fashion-specific and non-generic terms which are sufficiently detailed but not extremely lengthy to be exploited for constraining the generation. Motivated by recent findings in the field showing that humans tend to describe fashion items using only a few words [3], we propose to use noun chunks (*i.e.* short textual sentences composed of a noun along with its modifiers) that can effectively capture important information while reducing unnecessary words or details. Given that manually annotating all the images would be time-consuming and resource-intensive¹, we propose a novel framework to semi-automatically annotate the dataset using noun chunks. Firstly, domain-specific captions are collected from two available fashion datasets, namely FashionIQ [50] and Fashion200k [12], standardizing them with word lemmatization and eventually reducing each word to its root form with the NLTK library². Then, we extract noun chunks from the captions, filtering the results by removing all textual items that start with or contain special characters. After this pre-processing stage, we obtain more than 60k unique noun chunks, divided into three different categories (*i.e.* upper-body clothes, lower-body clothes, and dresses).

¹Since the Dress Code dataset consists of over 53k fashion items and assuming that each annotation requires approximately 5 minutes, a single annotator working 8 hours per day, 5 days a week, and 260 working days per year would take more than 2 years to complete the annotation task.

²<https://www.nltk.org/>

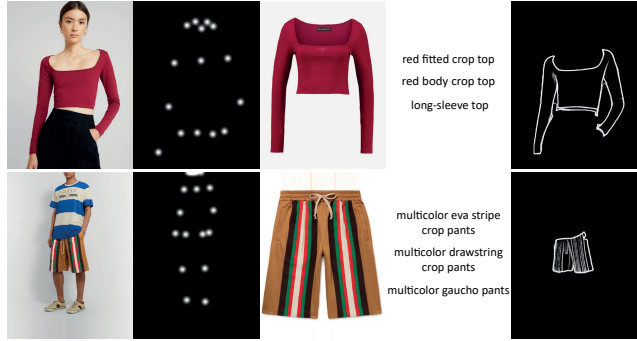


Figure 3. Sample images and multimodal data from our newly collected datasets.

Dataset	Text	Pose	Sketch	# Images	# Products	# Unique Texts	# Unique Words
VITON-HD [7]	✗	✓	✗	27,358	13,679	-	-
Dress Code [28]	✗	✓	✗	107,584	53,792	-	-
Be Your Own Prada [56]	✓	✓	✗	78,979	N/A	3,972	445
DF-Multimodal [19]	✓	✓	✗	44,096	N/A	10,253	77
VITON-HD Multimodal	✓	✓	✓	27,358	13,679	5,143	1,613
Dress Code Multimodal	✓	✓	✓	107,584	53,792	25,596	2,995

Table 1. Comparison of Dress Code and VITON-HD Multimodal with other fashion datasets with multimodal annotations.

To determine the most relevant noun chunks for each garment, we employ the CLIP model [34] and its open-source adaptation (*i.e.* OpenCLIP [49]). We select the ViT-L14@336 and RN50×64 models for CLIP, and the ViT-L14, ViT-H14, and ViT-g14 models for OpenCLIP. Prompt ensembling is performed to improve the results and, for each image, we select 25 noun chunks based on the top-5 noun chunks per model rated by cosine similarity between image and text embeddings, avoiding repetitions.

Fine-Grained Textual Annotation. To ensure the accuracy and representativeness of our annotations, we manually annotate a significant portion of Dress Code images. In particular, we select the three most representative noun chunks, among the 25 automatically associated, with each garment image. To minimize the annotation time, we develop a custom annotation tool that constrains the annotation time to an average time of 60 seconds per item and allows the annotator to manually insert noun chunks in the case that none of the automatically extracted ones are suitable for the image. Overall, we manually annotate 26,400 different garments (8,800 for each category) out of the 53,792 products included in the dataset, ensuring to include all fashion items of the original test set [28].

Coarse-Grained Textual Annotation. To complete the annotation, we first finetune the OpenCLIP ViT-B32 model, pre-trained on the English portion of the LAION5B dataset [40], using the newly annotated image-text pairs. We then use this model and the collected set of noun chunks to automatically tag all the remaining elements of the Dress Code dataset with the three most similar noun chunks, al-

ways determined via cosine similarity between multimodal embeddings. We employ the same strategy also to automatically annotate all garment images of the VITON-HD dataset. In this case, since this dataset only contains upper-body clothes, we limit the table noun chunks to the ones describing upper-body garments.

Extracting Sketches. The introduction of garment sketches can provide valuable design details that are not easily discernible from text alone. In this way, the dataset can provide a more accurate and comprehensive representation of the garments, leading to improved quality and better control of the generated design details. To extract sketches for both Dress Code and VITON-HD datasets, we employ PiDiNet [43], a pre-trained edge detection network.

Given that the selected datasets have originally been introduced for virtual try-on, they consist of both paired and unpaired test sets. While for the paired set we can directly use the human parsing mask to extract the garment of interest worn by the model and then feed it to the edge detection network, for the unpaired set we need to first create a warped version of the in-shop garment matching the body pose and shape of the target model. Following virtual try-on methods [47, 52], we train a geometric transformation module that performs a thin-plate spline transformation [36] of the input garment and then refines the warped result using a U-Net model [38]. From each warped garment, we extract the sketch image enabling the use of the proposed solution even in unpaired settings.

4.2. Comparison with Other Datasets

The only two text-to-image generation datasets available in the fashion domain [19, 56] are both based on images from the DeepFashion dataset [22]. While the dataset introduced in [56] contains short textual descriptions, DeepFashion-Multimodal [19] is annotated with attributes (*e.g.* category, color, fabric, etc.) that can be composed in longer captions. In Table 1, we summarize the main statistics of the publicly available datasets textual annotations compared with those of our newly extended datasets. As can be seen, our datasets contain more variety in terms of textual items and words, confirming the appropriateness of our annotation procedure and enabling a more personalized control of the generation process. Also, it is worth noting that the other datasets have no in-shop garment images making them difficult to employ in our case.

5. Experimental Evaluation

5.1. Implementation Details and Competitors

Training and Inference. All models are trained on the original splits of the Dress Code Multimodal and VITON-HD Multimodal datasets on a single NVIDIA A100 GPU for 150k steps, using a batch size of 16, a learning rate

of 10^{-5} with a linear warmup for the first 500 iterations, and AdamW [23] as optimizer with weight decay 10^{-2} . To speed up training and save memory, we use mixed precision [26]. We set both the fraction of steps conditioned by the sketch and the portion of masked conditions during training to 0.2. During inference, we employ the DDIM [42] with 50 steps as noise scheduler and set the classifier-free guidance parameter α to 7.5.

Baselines and Competitors. As first competitor, we use the out-of-the-box implementation of the inpainting Stable Diffusion pipeline³ provided by Huggingface. Moreover, we adapt two existing models, namely FICE [33] and SDEdit [25], to work on our setting. In particular, we re-train all main components of the FICE model on the newly collected datasets. We employ the same resolution used by the authors (*i.e.* 256×256), downsampling each image to 256×192 and applying padding to match the desired size (which is then removed during evaluation). To compare our model with a different conditioning strategy, we employ the approach proposed in [25] using our model trained only with text and human poses as input modalities and perform the sketch guidance using as starting latent variable the sketch image with added random noise. Following the original paper instructions, we use 0.8 as the strength parameter.

5.2. Evaluation Metrics

To assess the realism of generated images, we employ the Fréchet Inception Distance (FID) [16] and the Kernel Inception Distance (KID) [4]. For both metrics, we adopt the implementation proposed in [32]. Instead, to evaluate the adherence of the image to the textual conditioning input, we employ the CLIP Score (CLIP-S) [15] provided in the TorchMetrics library [9], using the OpenCLIP ViT-H/14 model as cross-modal architecture. We compute the score on the inpainted region of the generated output pasted on a 224×224 white background.

Pose Distance (PD). We propose a novel pose distance metric that measures the coherence of human body poses between the generated image and the original one estimating the distance between the human keypoints extracted from the original and generated images. Specifically, we employ the OpenPifPaf [20] human pose estimation network and compute the ℓ_2 distance between each pair of real-generated corresponding estimated keypoints. We only consider the keypoints involved in the generation (*i.e.* that falls in the mask M) and weigh each keypoint distance with the detector confidence to take into account any estimation errors.

Sketch Distance (SD). To quantify the adherence of the generated image to the sketch constraint, we propose a novel sketch distance metric. To compute the score, we extract the segmentation map of the original and generated

³<https://huggingface.co/runwayml/stable-diffusion-inpainting>

Model	Resolution	Modalities			Dress Code Multimodal					VITON-HD Multimodal				
		Text	Pose	Sketch	FID ↓	KID ↓	CLIP-S ↑	PD ↓	SD ↓	FID ↓	KID ↓	CLIP-S ↑	PD ↓	SD ↓
<i>Paired setting</i>														
Stable Diffusion [37]	256×192	✓			17.05	9.28	28.71	4.62	-	15.18	6.38	30.40	5.04	-
FICE [33]	256×192	✓	✓		30.63	23.54	28.72	6.87	-	49.44	44.74	29.26	6.37	-
MGD (ours)	256×192	✓	✓		5.57	1.67	31.33	2.37	-	10.11	3.14	31.85	2.90	-
<i>Paired setting</i>														
Stable Diffusion [37]	512×384	✓			17.43	9.48	29.18	9.24	0.467	16.28	6.56	30.70	10.78	0.410
SDEdit [25]	512×384	✓	✓	✓	10.19	5.03	29.21	5.41	0.398	13.07	4.66	30.58	6.76	0.306
MGD (ours)	512×384	✓	✓	✓	5.74	2.11	31.68	4.72	0.374	10.60	3.26	32.39	5.94	0.253
<i>Unpaired setting</i>														
Stable Diffusion [37]	256×192	✓			19.11	10.69	27.53	5.07	-	17.37	7.55	28.40	5.50	-
FICE [33]	256×192	✓	✓		34.14	26.86	26.03	7.15	-	52.74	48.58	25.94	6.58	-
MGD (ours)	256×192	✓	✓		7.01	2.19	29.58	2.96	-	11.54	3.18	29.95	3.30	-
<i>Unpaired setting</i>														
Stable Diffusion [37]	512×384	✓			19.55	10.80	28.02	9.89	0.582	18.45	7.87	28.74	11.60	0.561
SDEdit [25]	512×384	✓	✓	✓	11.38	5.69	27.10	6.16	0.509	15.12	5.67	28.61	7.35	0.406
MGD (ours)	512×384	✓	✓	✓	7.73	2.82	30.04	6.79	0.458	12.81	3.86	30.75	7.22	0.317

Table 2. Quantitative results on the Dress Code Multimodal and VITON-HD Multimodal datasets for both paired and unpaired settings .

Model	Modalities			Dress Code Multimodal				
	Text	Pose	Sketch	FID ↓	KID ↓	CLIP-S ↑	PD ↓	SD ↓
	✓			6.19	2.15	31.79	6.16	0.411
	✓	✓		6.31	2.33	31.67	5.31	0.405
MGD (ours)	✓	✓	✓	5.74	2.11	31.68	4.72	0.374
Model	Modalities			VITON-HD Multimodal				
	Text	Pose	Sketch	FID ↓	KID ↓	CLIP-S ↑	PD ↓	SD ↓
	✓			11.39	3.52	32.16	7.83	0.339
	✓	✓		11.07	3.36	32.27	6.77	0.318
MGD (ours)	✓	✓	✓	10.60	3.26	32.39	5.94	0.253

Table 3. Performance analysis on the paired setting of both datasets as input modalities vary.

garments using an off-the-shelf clothing segmentation network⁴. We then use the segmented garment area to extract garment sketches using the PIDInet [43] edge detector network. The final score is the mean squared error between these sketches, weighting the per-pixel results on the inverse pixel frequency of the activated pixels. More details about these proposed metrics can be found in the supplementary.

5.3. Experimental Results

Comparison with Other Methods. We test our proposal for the paired and unpaired settings of the considered datasets. In the former, the conditions (*e.g.* text, sketch) refers to the garment the model is wearing, while in the latter, the in-shop garment differs from the worn one. In Table 2, we report the quantitative results on Dress Code Multimodal and VITON-HD Multimodal in comparison with the aforementioned competitors. As can be seen, the proposed MGD model consistently outperforms competitors, in terms of realism (*i.e.* FID and KID) and coherency with input modalities (*i.e.* CLIP-S, PD, and SD). In particu-

⁴<https://github.com/levindabhi/cloth-segmentation>

	Modalities			Realism			Multimodal Coherence		
	Text	Pose	Sketch	Stable Diff.	FICE	SDEdit	Stable Diff.	FICE	SDEdit
Dress Code M.	✓			70.82	-	-	65.32	-	-
	✓	✓		70.73	96.26	-	65.15	84.48	-
	✓	✓	✓	70.29	-	52.54	65.38	-	66.23
VITON HD M.	✓			67.03	-	-	57.76	-	-
	✓	✓		66.17	93.84	-	73.73	83.46	-
	✓	✓	✓	60.71	-	53.44	69.47	-	59.34

Table 4. User study results on the unpaired setting of both datasets. We report the percentage of times an image from MGD is preferred against a competitor. Comparisons with FICE [33] are performed at 256 × 192 resolution.

lar, when considering low-resolution results, we notice that FICE [33] can produce images fairly consistent with the text conditioning, albeit less realistic than other methods. While Stable Diffusion [37] enhances image realism, it fails to preserve the input model’s pose due to the lack of pose information in the inputs. It is noteworthy that in this case, we compare the results of our model only using text and pose map as conditioning since both considered competitors are not conditioned on sketches. For this reason, we do not report the results in terms of sketch distance for low-resolution images.

In the high-resolution setting, we evaluate instead our MGD method using all multimodal conditions (*i.e.* text, pose map, and sketch) as input. In this case, we compare MGD with Stable Diffusion [37] plus SDEdit [25], where we use our text-pose conditioned denoising network as SDEdit backbone. Our findings indicate that Stable Diffusion performs worse in terms of the pose distance than both SDEdit and MGD, owing to the lack of pose information in the inputs. It is noteworthy that SDEdit performs worse than our model in all metrics. We attribute this behavior to the way sketch conditioning happens. In SDEdit,




Figure 4. Sample generated images on Dress Code Multimodal and VITON-HD Multimodal (bottom left) using all multimodal inputs.

it occurs only at the beginning by initializing z_t using the sketch image with added noise according to the conditioning strength, while our model conditions the denoising process in multiple steps, depending on the sketch conditioning parameter. Qualitative results reported in Fig. 4 highlight how our model better follows the given conditions and generate high-realistic images.

To validate our results based on human judgment, we conduct a user study that evaluates both the realism of the generation and the adherence to multimodal inputs. Overall we collect about 7k evaluations involving more than 150 users. Additional details are reported in the supplementary. Table 4 shows the user study results. Also in this case our model outperforms the competitors, thus confirming the effectiveness of our proposal.

Varying Input Modalities. In Table 3, we study the behavior of our MGD model when the input modalities are masked (*i.e.* where we feed the model with a zero tensor instead of the considered modality). In particular, we focus on the CLIP-S for text adherence and on the newly proposed pose and sketch distances for the pose and sketch coherency, respectively. Notice that the text input anchors the CLIP-S metrics of all experiments and makes them comparable in all cases. Starting from the fully conditioned model (*i.e.* text, pose, sketch), we mask the sketch. As the decrease of the sketch distance in Table 3 confirms, this input actually influences the generation process of our model in both the considered datasets. Also, this modality slightly affects the pose distance as the sketch implicitly contains information about the model’s body pose. We further mask the pose map input and compare the output with previous results. In

		Dress Code Multimodal					
Uncond. Portion	Sketch Cond.	FID ↓	KID ↓	CLIP-S ↑	PD ↓	SD ↓	
0.1	1.0	9.64	3.76	30.24	7.66	0.459	
0.2	1.0	8.62	3.24	29.06	7.51	0.430	
0.3	1.0	10.93	4.78	28.47	7.69	0.432	
0.2	0.8	8.56	3.28	29.31	7.32	0.433	
0.2	0.6	8.43	3.21	29.51	7.32	0.436	
0.2	0.4	8.11	3.00	29.79	7.13	0.440	
0.2	0.2	7.73	2.82	30.04	6.79	0.458	
0.2	0.0	7.82	2.85	29.93	6.26	0.519	

Table 5. Ablation analysis of our complete model varying the unconditioning portion during training and the sketch conditioning steps. Results refer to the unpaired setting.

this case, we can also notice a consistent difference with the text-only conditioned model, according to all metrics except CLIP-S as expected. These results confirm that our MGD model can effectively deal with the conditions in a disentangled way, making them optional.

Unconditional Training and Sketch Conditioning. In Table 5, we inquire about the fully conditioned network performance according to the variance of the portion of unconditional training. Additionally, we evaluate the results by varying the fraction of sketch conditioning steps. As can be seen, the best results are achieved by using 0.2 for both parameters. In particular, for unconditional training, we train three different models (*i.e.* with 0.1, 0.2, and 0.3). When evaluating the sketch conditioning parameter, we test our model with values between 0 and 1 with a stride of 0.2. It is worth noting that the sketch distance consistently decreases as the number of sketch conditioning steps increases, showing the robustness of the approach.

6. Conclusion

The Multimodal Garment Designer proposed in this paper is the first latent diffusion model defined for human-centric fashion image editing, conditioned by multimodal inputs such as text, body pose, and sketches. The novel architecture, trained on two new semi-automatically annotated datasets and evaluated with standard and newly proposed metrics, as well as by user studies, is very promising. The result is one of the first successful attempts to mimic the designers' job in the creative process of fashion design and could be a starting point for a capillary adoption of diffusion models in creative industries, oversight by human input.

Acknowledgments

This work has partially been supported by the European Commission under the PNRR-M4C2 project "FAIR - Future Artificial Intelligence Research" and the European Horizon 2020 Programme (grant number 101004545 - ReInHerit), and by the PRIN project "CREATIVE: CRoss-modal understanding and gENERATIOn of Visual and tEXtual content" (CUP B87G22000460001), co-funded by the Italian Ministry of University.

References

- [1] Omri Avrahami, Thomas Hayes, Oran Gafni, Sonal Gupta, Yaniv Taigman, Devi Parikh, Dani Lischinski, Ohad Fried, and Xi Yin. SpaText: Spatio-Textual Representation for Controllable Image Generation. In *CVPR*, 2023. 4
- [2] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. eDiff-I: Text-to-Image Diffusion Models with an Ensemble of Expert Denoisers. *arXiv preprint arXiv:2211.01324*, 2022. 4
- [3] Federico Bianchi, Jacopo Tagliabue, and Bingqing Yu. Query2Prod2Vec: Grounded word embeddings for eCommerce. In *NAACL*, 2021. 5
- [4] Mikolaj Bińkowski, Dougal J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. In *ICLR*, 2018. 6
- [5] Shin-I Cheng, Yu-Jie Chen, Wei-Chen Chiu, Hung-Yu Tseng, and Hsin-Ying Lee. Adaptively-Realistic Image Generation from Stroke and Sketch with Diffusion Model. In *WACV*, 2023. 2
- [6] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. In *ICCV*, 2021. 2
- [7] Seunghwan Choi, Sunghyun Park, Minsoo Lee, and Jaegul Choo. VITON-HD: High-Resolution Virtual Try-On via Misalignment-Aware Normalization. In *CVPR*, 2021. 1, 2, 5
- [8] Guillem Cucurull, Perouz Taslakian, and David Vazquez. Context-aware visual compatibility prediction. In *CVPR*, 2019. 1
- [9] Nicki Skaftø Detlefsen, Jiri Borovec, Justus Schock, Ananya Harsh Jha, Teddy Koker, Luca Di Liello, Daniel Stancl, Changsheng Quan, Maxim Grechkin, and William Falcon. TorchMetrics-Measuring Reproducibility in PyTorch. *Journal of Open Source Software*, 7(70):4101, 2022. 6
- [10] Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In *NeurIPS*, 2021. 2
- [11] M Hadi Kiapour, Xufeng Han, Svetlana Lazebnik, Alexander C Berg, and Tamara L Berg. Where to buy it: Matching street clothing photos in online shops. In *ICCV*, 2015. 1
- [12] Xintong Han, Zuxuan Wu, Phoenix X Huang, Xiao Zhang, Menglong Zhu, Yuan Li, Yang Zhao, and Larry S Davis. Automatic spatially-aware fashion concept discovery. In *ICCV*, 2017. 5
- [13] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S Davis. Viton: An image-based virtual try-on network. In *CVPR*, 2018. 1
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 2015. 4
- [15] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. In *EMNLP*, 2021. 6
- [16] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, Günter Klambauer, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a Nash equilibrium. In *NeurIPS*, 2017. 6
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *NeurIPS*, 2020. 2
- [18] Wei-Lin Hsiao and Kristen Grauman. Creating capsule wardrobes from fashion images. In *CVPR*, 2018. 1
- [19] Yuming Jiang, Shuai Yang, Haonan Qju, Wayne Wu, Chen Change Loy, and Ziwei Liu. Text2human: Text-driven controllable human image generation. *ACM Transactions on Graphics*, 41(4):1–11, 2022. 2, 5, 6
- [20] Sven Kreiss, Lorenzo Bertoni, and Alexandre Alahi. OpenPifPaf: Composite Fields for Semantic Keypoint Detection and Spatio-Temporal Association. *IEEE Transactions on Intelligent Transportation Systems*, 23(8):13498–13511, 2021. 6
- [21] Wenbo Li, Zhe Lin, Kun Zhou, Lu Qi, Yi Wang, and Jiaya Jia. MAT: Mask-Aware Transformer for Large Hole Image Inpainting. In *CVPR*, 2022. 4
- [22] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations. In *CVPR*, 2016. 1, 6
- [23] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *ICLR*, 2019. 6
- [24] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. RePaint: Inpainting Using Denoising Diffusion Probabilistic Models. In *CVPR*, 2022. 4
- [25] Chenlin Meng, Yutong He and Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *ICLR*, 2022. 2, 6, 7
- [26] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael

- Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. Mixed Precision Training. In *ICLR*, 2018. 6
- [27] Davide Morelli, Alberto Baldrati, Giuseppe Cartella, Marcella Cornia, Marco Bertini, and Rita Cucchiara. LaDIVTON: Latent Diffusion Textual-Inversion Enhanced Virtual Try-On. In *ACM Multimedia*, 2023. 1
- [28] Davide Morelli, Matteo Fincato, Marcella Cornia, Federico Landi, Fabio Cesari, and Rita Cucchiara. Dress Code: High-Resolution Multi-Category Virtual Try-On. In *ECCV*, 2022. 1, 2, 5
- [29] Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhong-gang Qi, Ying Shan, and Xiaoju Qie. T2I-Adapter: Learning Adapters to Dig out More Controllable Ability for Text-to-Image Diffusion Models. *arXiv preprint arXiv:2302.08453*, 2023. 2, 3
- [30] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, 2021. 2
- [31] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: Towards Photorealistic Image Generation and Editing with Text-Guided Diffusion Models. In *ICML*, 2022. 2
- [32] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On Aliased Resizing and Surprising Subtleties in GAN Evaluation. In *CVPR*, 2022. 6
- [33] Martin Pernuš, Clinton Fookes, Vitomir Štruc, and Simon Dobrišek. FICE: Text-Conditioned Fashion Image Editing With Guided GAN Inversion. *arXiv preprint arXiv:2301.02110*, 2023. 2, 6, 7
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, A. Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, 2021. 2, 5
- [35] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical Text-Conditional Image Generation with CLIP Latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [36] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *CVPR*, 2017. 6
- [37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis With Latent Diffusion Models. In *CVPR*, 2022. 2, 3, 7
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI*, 2015. 6
- [39] Rohan Sarkar, Navaneeth Bodla, Mariya I Vasileva, Yen-Liang Lin, Anurag Beniwal, Alan Lu, and Gerard Medioni. OutfitTransformer: Learning Outfit Representations for Fashion Recommendation. In *WACV*, 2023. 1
- [40] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022. 5
- [41] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 2, 3
- [42] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. In *ICLR*, 2021. 6
- [43] Zhuo Su, Wenzhe Liu, Zitong Yu, Dewen Hu, Qing Liao, Qi Tian, Matti Pietikäinen, and Li Liu. Pixel difference networks for efficient edge detection. In *ICCV*, 2021. 6, 7
- [44] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *WACV*, 2022. 4
- [45] Ming Tao, Hao Tang, Fei Wu, Xiao-Yuan Jing, Bing-Kun Bao, and Changsheng Xu. DF-GAN: A Simple and Effective Baseline for Text-to-Image Synthesis. In *CVPR*, 2022. 2
- [46] Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. Sketch-Guided Text-to-Image Diffusion Models. *arXiv preprint arXiv:2211.13752*, 2022. 2
- [47] Bochao Wang, Huabin Zheng, Xiaodan Liang, Yimin Chen, Liang Lin, and Meng Yang. Toward characteristic-preserving image-based virtual try-on network. In *ECCV*, 2018. 1, 6
- [48] Tengfei Wang, Ting Zhang, Bo Zhang, Hao Ouyang, Dong Chen, Qifeng Chen, and Fang Wen. Pretraining is All You Need for Image-to-Image Translation. *arXiv preprint arXiv:2205.12952*, 2022. 2
- [49] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *CVPR*, 2022. 5
- [50] Hui Wu, Yupeng Gao, Xiaoxiao Guo, Ziad Al-Halah, Steven Rennie, Kristen Grauman, and Rogerio Feris. Fashion IQ: A New Dataset Towards Retrieving Images by Natural Language Feedback. In *CVPR*, 2021. 5
- [51] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. AttnGAN: Fine-Grained Text to Image Generation with Attentional Generative Adversarial Networks. In *CVPR*, 2018. 2
- [52] Han Yang, Ruimao Zhang, Xiaobao Guo, Wei Liu, Wang-meng Zuo, and Ping Luo. Towards photo-realistic virtual try-on by adaptively generating-preserving image content. In *CVPR*, 2020. 1, 6
- [53] Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. Cross-modal contrastive learning for text-to-image generation. In *CVPR*, 2021. 2
- [54] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023. 2
- [55] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. DM-GAN: Dynamic Memory Generative Adversarial Networks for Text-to-Image Synthesis. In *CVPR*, 2019. 2
- [56] Shizhan Zhu, Raquel Urtasun, Sanja Fidler, Dahua Lin, and Chen Change Loy. Be Your Own Prada: Fashion Synthesis with Structural Coherence. In *ICCV*, 2017. 2, 5, 6