

Be Your Own Prada: Fashion Synthesis with Structural Coherence

Shizhan Zhu¹ Sanja Fidler^{2,3} Raquel Urtasun^{2,3,4}
Dahua Lin¹ Chen Change Loy¹

¹Department of Information Engineering, The Chinese University of Hong Kong

²University of Toronto, ³Vector Institute, ⁴Uber Advanced Technologies Group

{szzhu, dhlin, cclloy}@ie.cuhk.edu.hk, {fidler, urtasun}@cs.toronto.edu

Abstract

We present a novel and effective approach for generating new clothing on a wearer through generative adversarial learning. Given an input image of a person and a sentence describing a different outfit, our model “redresses” the person as desired, while at the same time keeping the wearer and her/his pose unchanged. Generating new outfits with precise regions conforming to a language description while retaining wearer’s body structure is a new challenging task. Existing generative adversarial networks are not ideal in ensuring global coherence of structure given both the input photograph and language description as conditions. We address this challenge by decomposing the complex generative process into two conditional stages. In the first stage, we generate a plausible semantic segmentation map that obeys the wearer’s pose as a latent spatial arrangement. An effective spatial constraint is formulated to guide the generation of this semantic segmentation map. In the second stage, a generative model with a newly proposed compositional mapping layer is used to render the final image with precise regions and textures conditioned on this map. We extended the DeepFashion dataset [8] by collecting sentence descriptions for 79K images. We demonstrate the effectiveness of our approach through both quantitative and qualitative evaluations. A user study is also conducted.

1. Introduction

Imagine that you could be your own fashion designer, and be able to seamlessly transform your current outfit in the photo into a completely new one, by simply describing it in words (Figure 1). In just minutes you could design and “try on” hundreds of different shirts, dresses, or even styles, allowing you to easily discover what you look good in. The goal of this paper is to develop a method that can generate new outfits onto existing photos, in a way that preserves *structural coherence* from multiple perspectives:

1. Retaining the body shape and pose of the wearer,



Figure 1. Given an original wearer’s input photo (left) and different textual descriptions (second column), our model generates new outfits onto the photograph (right three columns) while preserving the pose and body shape of the wearer.

2. Producing regions and the associated textures that conform to the language description, and
3. Enforcing coherent visibility of body parts.

Meeting *all* these requirements at the same time is a very challenging task. First, the input image is the only source from which we can mine for the body shape information. With only a single view of the wearer, it is nontrivial to recover the body shape accurately. Moreover, we do not want the shape of the generated outfit to be limited by the original garments of the wearer. For example, replacing the original long-sleeve shirt with a short-sleeve garment would require the model to hallucinate the person’s arms and skin.

Conventional 2D non-parametric methods or 3D graphics approaches meet the first requirement through structural constraints derived from human priors. They can be in the form of accurate physical measurements (*e.g.* height, waist, hip, arm length) to create 3D virtual bodies [10]; manual manipulations of sliders such as height, weight and waist girth [19]; or indication of joint positions and a rough sketch outlining the human body silhouette [16]. All these methods require explicit human interventions at test time,

which would limit their applicability in practical settings. In addition, as these methods provide no obvious ways to incorporate textual descriptions to condition the synthesis process, it is non-trivial to fulfil the second requirement with existing methods. Lastly, they do not meet the third requirement as they do not support hallucination of the missing parts.

Generative Adversarial Networks (GAN) [2] is an appealing alternative to conventional methods. In previous work, DCGAN [11], a GAN formulation combined with convolutional networks, has been shown to be an effective model to produce realistic images. Moreover, it allows for an end-to-end embedding of textual descriptions to condition the image generation. The task of clothing generation presents two significant challenges which are difficult to address with the standard DCGAN. First, it directly targets the pixel values and provides no mechanism to enforce structural coherence w.r.t. to the input. Second, it tends to average out the pixels [14], thus resulting in various artifacts, *e.g.* blurry boundaries, as shown by our experiments.

To tackle this problem, we propose an effective two-stage GAN framework that generates shape and textures in different stages. In the first stage, we aim to generate a plausible human segmentation map that specifies the regions for body parts and the upper-body garment. This stage is responsible for preserving the body shape and ensuring the coherent visibility of parts based on the description. In the second stage, the generator takes both the produced segmentation map and the textual description as conditions, and renders the region-specific texture onto the photograph.

To ensure the coherence in structure of the synthesized image with respect to the input image (*i.e.* preserving the body shape and pose of the wearer), we present an effective spatial constraint that can be derived from the input photograph. We formulate it carefully so that it does not contradict to the textual description when both of them are used to condition the first-stage GAN. In addition, we also introduce a new compositional mapping layer into the second-stage GAN to enforce region-specific texture rendering guided by the segmentation map. In contrast to existing GANs that perform non-compositional synthesis, the new mapping layer is capable of generating more coherent visibility of body parts with image region-specific textures.

To train our model, we extend the DeepFashion dataset [8] by annotating a subset of 79K upper-body images with sentence descriptions and human body annotations¹. Extensive quantitative and qualitative comparisons are performed against existing GAN baselines and 2D non-parametric approaches. We also conduct a user study in order to obtain an objective evaluation on both the shape and image generation results.

¹The data and code can be found at <http://mmlab.ie.cuhk.edu.hk/projects/FashionGAN/>.

2. Related Work

Generative Adversarial Networks (GAN) [2] have shown impressive results generating new images, *e.g.* faces [11], indoor scenes [15], fine-grained objects like birds [12], or clothes [17]. Training GANs based on conditions incorporates further information to guide the generation process. Existing works have explored various conditions, from category labels [9], text [12] to an encoded feature vector [17]. Different from the studies above, our study aims at generating the target by using the spatial configuration of the input images as a condition. The spatial configuration is carefully formulated so that it is agnostic to the clothing worn in the original image, and only captures information about the user's body.

There exist several studies to transfer an input image to a new one. Ledig *et al.* [6] apply the GAN framework to super-resolve a low-resolution image. Zhu *et al.* [3] use a conditional GAN to transfer across the image domains, *e.g.* from edge maps to real images, or from daytime images to night-time. Isola *et al.* [3] change the viewing angle of an existing object. Johnson *et al.* [4] apply GANs to neural style transfer. All these studies share a common feature - the image is transformed *globally* on the texture level but is not region-specific. In this study, we explore a new compositional mapping method that allows region-specific texture generation, which provides richer textures for different body regions.

There are several recent studies that explore improved image generation by stacking GANs. Our work is somewhat similar in spirit to [15, 18] – our idea is to have the first stage to create the basic composition, and the second stage to add the necessary refinements to the image generated in the first stage. However, the proposed FashionGAN differs from S²GAN [15] in that the latter aims at synthesizing a surface map from a random vector in its first stage. In contrast, our goal is to generate a plausible mask whose structure conforms to a given photograph and language description, which requires us to design additional spatial constraints and design coding as conditions. Furthermore, these two conditions should not contradict themselves. Similarly, our work requires additional constraints which are not explored in [18]. Compositional mapping is not explored in the aforementioned studies as well.

Yo *et al.* [17] propose an image-conditional image generation model to perform domain transfer, *e.g.*, generating a piece of clothing from an input image of a dressed person. Our work differs in that we aim at changing the outfit of a person into a newly designed one based on a textual description. Rendering new outfits onto photographs with unconstrained human poses bring additional difficulties in comparison with work that generates pieces of clothing in a fixed view-angle as in [17].

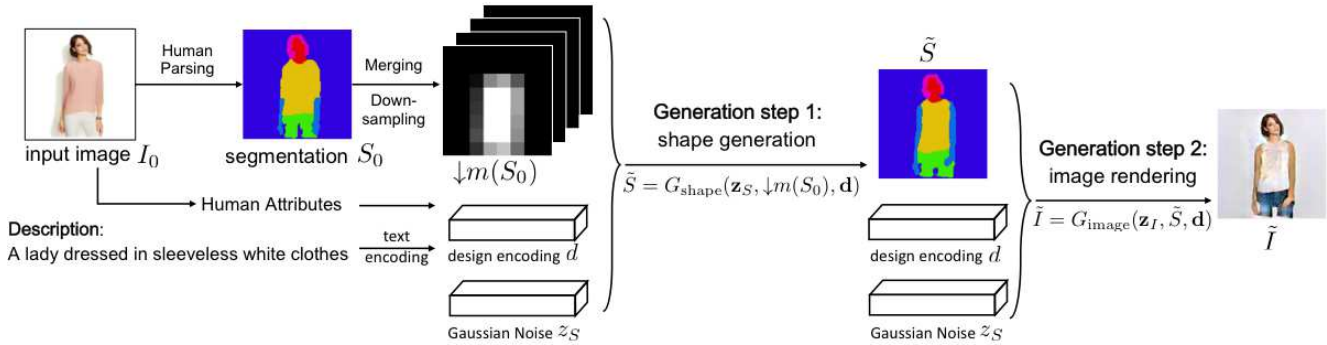


Figure 2. **Proposed framework.** Given an input photograph of a person and a sentence description of a new desired outfit, our model first generates a segmentation map \tilde{S} using the generator from the first GAN. We then render the new image with another GAN, with the guidance from the segmentation map generated in the previous step. At test time, we obtain the final rendered image with a forward pass through the two GAN networks.

3. Methodology

Our framework is inspired by the generative adversarial network (GAN) proposed by Goodfellow *et al.* [2]. We first provide a concise review of GAN, and then introduce our outfit generation framework. Generative Adversarial Network [2] has shown a powerful capability in generating realistic natural images. A typical GAN contains a generator G and a discriminator D . They are jointly trained with a learning objective given below:

$$\min_G \max_D \mathbb{E}_{I \sim p_{\text{data}}} [\log D(I)] + \mathbb{E}_{\mathbf{z} \sim p_z} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

Here, \mathbf{z} is a random or encoded vector, p_{data} is the empirical distribution of training images, and p_z is the prior distribution of \mathbf{z} . It was proven in [2] that when it reaches the maximum, the distribution of $G(\mathbf{z})$ would converge to p_{data} , where the discriminator cannot distinguish the images $I \sim p_{\text{data}}$ from the generated ones.

3.1. Overview of FashionGAN

We define the problem as follows. We assume we have the original image of a wearer and a sentence description of the new outfit. An example of a description we envision is “a white blouse with long sleeves but without a collar”. Our goal is to produce a new image of the user wearing the desired outfit.

Our method requires training data in order to learn the mapping from one photo to the other given the description. We do not assume paired data where the same user is required to wear two outfits (current, and the described target outfit). Instead, we only require one photo per user where each photo has a sentence description of the outfit. Such data is much easier to collect (Sec. 3.5).

Since in our scenario we only have one (described) image per user, this image serves as both the input and the target during training. Thus, rather than working directly

with the original image I_0 , we extract the person’s segmentation map, S_0 , which contains pixel-wise class labels such as *hair, face, upper-clothes, pants/shorts*, etc. The segmentation map is thus capturing the shape of the wearer’s body and parts, but not their appearance.

To capture further information about the wearer, we extract a vector of binary attributes, \mathbf{a} , from the person’s face, body and other physical characteristics. Examples of attributes include gender, long/short hair, wearing/not wearing sunglasses and wearing/not wearing hat. The attribute vector may additionally capture the mean RGB values of skin color, as well as the aspect ratio of the person, representing coarse body size. These are the properties that our final generated image should ideally preserve. Details of how we extract this information are given in Sec. 3.5.

We represent the description as a vector \mathbf{v} using an existing text encoder (details in Sec. 3.5). Our problem is then formalized as follows. Given $\mathbf{d} = (\mathbf{a}, \mathbf{v})$, which we call the *design coding*, and the human segmentation map S_0 , our goal is to synthesize a new high-quality image \tilde{I} of the wearer matching the requirements provided in the description, while at the same time preserving the wearer’s pose and body shape. Note that during training, $\tilde{I} = I_0$.

As shown in Fig. 2, we decompose the overall generative process into two relatively easier stages, namely the human segmentation (shape) generation (corresponding to the desired/target outfit) and texture rendering. This decomposition can be expressed as follows:

$$\tilde{S} \leftarrow G_{\text{shape}}(\mathbf{z}_S, \downarrow m(S_0), \mathbf{d}), \quad (2)$$

$$\tilde{I} \leftarrow G_{\text{image}}(\mathbf{z}_I, \tilde{S}, \mathbf{d}). \quad (3)$$

Here, G_{shape} and G_{image} are two separate generators.

More precisely, in our first stage (Eq. (2)), we first generate a human segmentation map \tilde{S} by taking the original segmentation map S_0 and the design coding \mathbf{d} into account. Here $\downarrow m(S_0)$ is a low-resolution representation of S_0 , serving as the **spatial constraint** to ensure structural coherence



Figure 3. This figure motivates the use of $\downarrow m(S_0)$ as a spatial constraint in the first-stage of FashionGAN. Using the high-resolution segmentation map S_0 will result in artifacts in the final generated image, when the segmentation maps convey semantic meanings that differ from the textual description. See Sec. 3.2 for details.

of the generated map \tilde{S} to the body shape and pose of the wearer. In the second stage (Eq. (3)), we use the generated segmentation map \tilde{S} produced by the first generator, as well as the design coding \mathbf{d} , to render the garments for redressing the wearer. The texture for each semantic part is generated in different specialized channels, which are then combined according to the segmentation map \tilde{S} to form the final rendered image. We call this process a **compositional mapping**. This newly introduced mapping is useful for generating high-quality texture details within specific regions.

We provide details of the two generators in our framework in Sec. 3.2 and Sec. 3.3, respectively.

3.2. Segmentation Map Generation (G_{shape})

Our first generator G_{shape} aims to generate the semantic segmentation map \tilde{S} by conditioning on the spatial constraint $\downarrow m(S_0)$, the design coding $\mathbf{d} \in \mathbb{R}^D$, and the Gaussian noise $\mathbf{z}_S \in \mathbb{R}^{100}$. We now provide more details about this model. To be specific, assume that the original image is of height m and width n , i.e., $I_0 \in \mathbb{R}^{m \times n \times 3}$. We represent the segmentation map S_0 of the original image using a pixel-wise one-hot encoding, i.e., $S_0 \in \{0, 1\}^{m \times n \times L}$, where L is the total number of labels. In our implementation, we use $L = 7$ corresponding to *background*, *hair*, *face*, *upper-clothes*, *pants/shorts*, *legs*, and *arms*.

Spatial Constraint $\downarrow m(S_0)$. We merge and down-sample the original segmentation map S_0 into $\downarrow m(S_0) \in [0, 1]^{m' \times n' \times L'}$ ($L' < L$ in our implementation), as a conditioning variable to G_{shape} . In particular, we use $L' = 4$ categories: *background*, *hair*, *face*, and *rest*. This essentially maps all the clothing pixels into a generic *rest* (or *body*) class. Thus, $\downarrow m(S_0)$ is agnostic of the clothing worn in the original image, and only captures information about the user’s body. This spatial constraint plays an important role in preserving structural coherence of the generated shape \tilde{S} , while still allowing variability in the generative process.

We use a down-sampled version of S_0 as a constraint so as to weaken the correlation between the two conditions S_0 and \mathbf{d} , which can contradict each other. Specifically, while S_0 keeps the complete information of the wearer’s body shape, its internal partitioning of regions do not necessarily agree with the specifications conveyed in the design




Figure 4. Conditioning on the same input image, the G_{shape} generates different human segmentation maps based on different design coding \mathbf{d} and a random vector \mathbf{z} . We can observe clear shape differences in sleeves, the length of the upper-clothes, and the labels assigned to legs/pants across the different samples.

coding \mathbf{d} . If we were to directly feed the high-resolution segmentation map of the original image into the model, strong artifacts would appear when the textual description contradicts with the segmentation map, e.g., the model simultaneously receives the text description “to generate a long dress” and the segmentation map that indicates short upper clothes. Figure 3 shows such failure cases.

Shape Generation. We want our G_{shape} to output a new human segmentation map $\tilde{S} \in [0, 1]^{m \times n \times L}$. This output should ideally have attributes consistent with the design coding \mathbf{d} , while at the same time, the generated human shape should conform to the human pose as encoded in the original S_0 . The generated segmentation map \tilde{S} should differ from the original human shape S_0 with new variations introduced by the design coding \mathbf{d} and noise \mathbf{z}_S . Figure 4 illustrates an example of the generated segmentation map. We observe that while the length of the sleeve and upper-clothes vary in different generated samples, the human pose and body shape remain consistent.

To produce the segmentation map \tilde{S} , we employ a GAN to learn the generator G_{shape} . Both the generator and discriminator comprise of convolution / deconvolution layers with batch normalization and non-linear operations. Note that different from most of the existing GANs for image generation, the shape map \tilde{S} we are generating in this step is governed by additional constraints – each pixel in the map has a probabilistic simplex constraint, i.e. $\tilde{S}_{ij} \in \Delta^L, 1 \leq i \leq m, 1 \leq j \leq n$. We use the Softmax activation function on each pixel at the end of the generator, so that the generated fake shape map is comparable with the real segmentation map. We observe that the GAN framework can also learn well in this scenario. Please refer to *suppl. material* for a detailed description of the network structure.

3.3. Texture Rendering (G_{image})

Having obtained the human segmentation map $\tilde{S} \in \mathbb{R}^{m \times n \times L}$ from the generator G_{shape} , we now use this map along with the design coding vector $\mathbf{d} \in \mathbb{R}^D$ to render the final image $\tilde{I} \in \mathbb{R}^{m \times n \times 3}$ using the second-stage G_{image} .

Compositional Mapping. Conventional GANs generate an image without enforcing region-specific texture rendering.

In the proposed FashionGAN, we propose a new compositional mapping layer that generates the image with the guidance of the segmentation map. In comparison to non-compositional counterparts, the new mapping layer helps to generate textures more coherent to each region and maintain visibility of body parts.

Formally, we train a specific channel in G_{image} for each category l , where $1 \leq l \leq L$, and L is the total number of labels in the segmentation map. We denote the set of pixels that belong to category l as P_l , and form our final generated image \tilde{I} as a collection of pixels $(\tilde{I})_p$ indexed by p ,

$$(\tilde{I})_p = \sum_{l=1}^L \mathbf{1}_{p \in P_l} \cdot (\tilde{I}_l)_p, \quad (4)$$

where p is the index of the pixel and \tilde{I}_l is the specific channel for the l -th category. Here $\mathbf{1}_{(\cdot)}$ is an indicator function.

Image Generation. Similar to the networks in Sec. 3.2, the generator and discriminator in this step are also composed of convolution / deconvolution layers with batch normalization and non-linear operations. Instead of assigning a Tanh activation function at the end of the network as most of GAN architectures do, we put this activation before the region-specific rendering layer \tilde{I}_l . This is important for achieving a stable combination of all the channels generated by the network. Please refer to *supplementary material* for a detailed description of the network structure.

3.4. Training

Our two GANs are trained separately due to the non-differentiable argmax operation between the two steps. The training process needs one fashion image I_0 for each person in the training set, along with the textual description (represented by its designing coding \mathbf{d}) and the segmentation map S_0 . In our first GAN, we derive the tuple $\{\downarrow m(S_0), \mathbf{d}, S_0\}$ from each training sample and train the networks, following the typical conditional GAN training procedure. In our second GAN, we derive the tuple $\{\downarrow m(S_0), \mathbf{d}, I_0\}$ from each training sample for training. We use the Adam optimizer [5] in training. Discriminative networks only appear in the training phase. Similar to [12], we provide the conditions (design coding, segmentation maps) to the discriminative networks to enforce consistency between the conditions and the generated results.

3.5. Implementation Details and Dataset

The dimensionality of the design coding \mathbf{d} is $D = 50$. Ten dimensions in \mathbf{d} serve as the human attributes. We represent the binary attributes of gender, long/short hair, w/o sunglasses, w/o hat with one dimension each. We extract the medium value of the R, G, B as well as the Y (gray) channel among the skin region, a total of four dimensions, to represent the skin color. We use the height and width of the given person to represent the size as well as the

aspect ratio. The remaining 40 dimensions are the encoded text. We follow [12] to construct the text encoder, which can be jointly tuned in each of the GANs in our framework. The resolution of our output image \tilde{I} is 128×128 (i.e. $m = n = 128$).

We perform bicubic down-sampling to get $\downarrow m(S_0)$, with a size of 8×8 (i.e. $m' = n' = 8$). We keep the hair and face regions in our merged maps avoiding the need for the generator to generate the exact face as the original wearer (we replace the generated hair/face region with the original image I_0). It is hard and not necessary in practice.

To train our framework we extended the publicly available DeepFashion dataset [8] with richer annotations (captions and segmentation maps). In particular, we selected a subset of 78,979 images from the DeepFashion attribute dataset, in which the person is facing toward the camera, and the background of the image is not severely cluttered. Training our algorithm requires segmentation maps and captions for each image. We manually annotated one sentence per photo, describing only the visual facts (e.g., the color, texture of the clothes or the length of the sleeves), avoiding any subjective assessments. For segmentation, we first applied a semantic segmentation method (VGG model fine-tuned on the ATR dataset [7]) to all the images, and then manually checked correctness. We manually relabeled the incorrectly segmented samples with GrabCut [13].

4. Experiments

We verify the effectiveness of FashionGAN through both quantitative and qualitative evaluations. Given the subjective nature of fashion synthesis, we also conduct a blind user study to compare our method with 2D nonparametric based method and other GAN baselines.

Benchmark. We randomly split the whole dataset (78,979 images) into a disjoint training set (70,000 images) and test set (8,979 images). All the results shown in this section are drawn from the test set. A test sample is composed of a given (original) image and a sentence description serving as the redressing condition.

Baselines. As our problem requires the model to generate a new image by keeping the person’s pose, many existing unconditional GAN-based approaches (e.g., DCGAN [11]) are not directly applicable to our task. Instead, we use the conditional variants to serve as the baseline approach in our evaluation. We compare with several baselines as follows:

(1) One-step GAN: To demonstrate the effectiveness of the proposed two-step framework, we implemented a conditional GAN to directly generate the final image in one step, i.e., $\tilde{I} = G_{\text{direct}}(\mathbf{z}, S_0, \mathbf{d})$. We refer to this type of baseline as *One-Step*. Since we aim to generate a new outfit that is consistent with the wearer’s pose in the original photo, the one-step baseline also requires similar spatial priors. Recall that we need to avoid contradiction between

Table 1. Evaluation results of detecting five structure-relevant attributes from synthesized images. Average precision is reported.

Images from	Has T-Shirt	Has Long Sleeves	Has Shorts	Has Jeans	Has Long Pants	mAP
Original (Upper-bound)	77.6	88.2	90.4	86.5	91.2	86.8
One-Step-8-7	56.5	73.4	79.6	73.5	79.1	72.4
One-Step-8-4	57.1	75.0	80.1	74.3	79.8	73.3
Non-Compositional	54.7	79.2	82.3	72.8	84.4	74.7
Non-Deep	56.1	69.4	75.7	74.3	76.5	71.8
FashionGAN	63.2	86.9	90.0	82.1	90.7	82.6

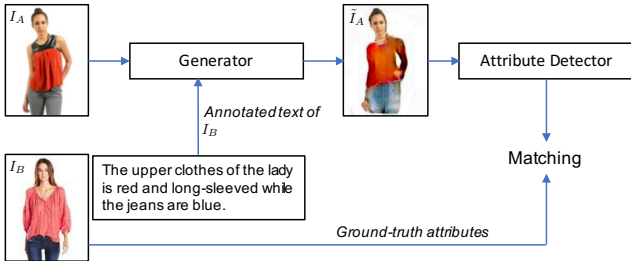


Figure 5. The experimental procedure to quantitatively verify the attribute and structure consistency of a generated image.

conditions from the text description and segmentation (see Sec. 3.1). Hence, for a fair comparison between our proposed approach and this baseline, we feed in the down-sampled version of the ground-truth segmentation map. We further divide this type of baseline into two different settings based on the way we use the shape prior S_0 :

- *One-Step-8-7*: We use the down-sampled but not merged segmentation map ($8 \times 8 \times 7$) as the prior;
- *One-Step-8-4*: We use the down-sampled merged segmentation map ($8 \times 8 \times 4$) as the prior (the same setting we used in our first stage GAN).

The architecture of the generator and discriminator used in these baselines are consistent to those used in our proposed method, i.e., containing 6 deconvolution and convolution layers in both the generator and discriminator.

(2) **Non-Compositional**: To demonstrate the effectiveness of the segmentation guidance, we build a baseline that generates an image as a whole, i.e., without using Eq. (4). In this baseline, we use two generative stages as in our proposed framework. In addition, the first stage generator of this baseline is still conditioned on the spatial constraint to ensure structure coherence to the wearer’s pose.

4.1. Quantitative Evaluation

A well-generated fashion image should faithfully produce regions and the associated textures that conform to the language description. This requirement can be assessed by examining if the desired outfit attributes are well captured by the generated image. In this section, we conduct a quantitative evaluation of our approach to verify the capability of FashionGAN in preserving attribute and structural coherence with the input text.

We selected a few representative attributes from DeepFashion, namely, ‘Has T-Shirt’, ‘Has Long Sleeves’, ‘Has Shorts’, ‘Has Jeans’, ‘Has Long Pants’. These attributes are all *structure-relevant*. A generative model that is poor in maintaining structural coherence will perform poorly on these attributes. Specifically, we performed the following experiment, as illustrated in Fig. 5. (1) For each test image I_A , we used a sentence of another randomly selected image I_B as the text input. The same image-text pairs were kept for all baselines for a fair comparison. (2) We used the image-text pair as input and generated a new image \tilde{I}_A using a generative model. (3) We used an external attribute detector² to predict the attributes on \tilde{I}_A . (4) Attribute prediction accuracy was computed by verifying the predictions on \tilde{I}_A against ground-truth attributes on I_B .

Table 1 summarizes the attribute prediction results. It can be observed that attribute predictions yielded by FashionGAN are more accurate than other baselines. In particular, our approach outperforms one-step GANs that come without the intermediate shape generation, and two-stage GAN that does not perform compositional mapping. Moreover, the performance of FashionGAN is close to the upper-bound, which was provided by applying the attribute detector on image I_B , where the text input originated from. The results suggest the superiority of FashionGAN in generating fashion images with structure coherence.

4.2. Qualitative Evaluation

Conditioning on the Same Wearer. Given an image, we visualize the output of FashionGAN with different sentence descriptions. We show all the intermediate results and final rendering step-by-step in Fig. 6, showcasing our generation process. A plausible segmentation map is generated first, and one can notice the variation in shape (e.g., the length of the sleeve). The image generated in the second step has consistent shape with the shape generated in the first step. The generated samples demonstrate variations in textures and colors, while the body shape and pose of the wearer are retained.

Conditioning on the Same Description. In this experiment, we choose photos of different wearers but use the same description to redress them. We provide results in

²We fine-tuned the R*CNN model [1] on our training set to serve as our attribute detector.

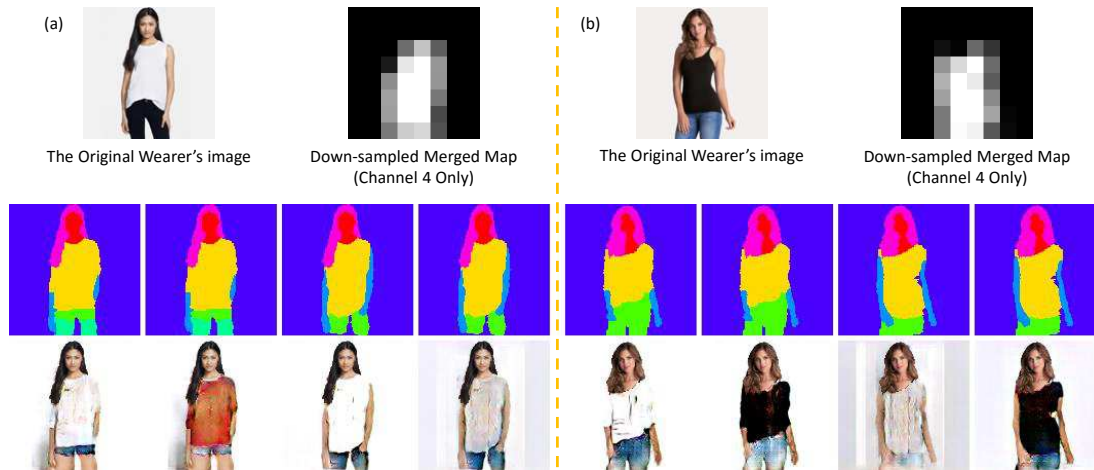


Figure 6. **Conditioning on the same wearer.** A step-by-step visualization of the generation process on two original images (here the sentence descriptions are different for all the 8 samples). Based on the original image and the down-sampled merged segmentation map (top), we first generate a plausible human shape (the middle row), and then use this as priors to guide the texture synthesis in the succeeding step (the bottom row).



Figure 7. **Conditioning on the same sentence description.** Samples conditioned on the same sentence “A lady dressed in blue clothes”.

Fig. 7. Regardless of the variations in the human body shapes and poses, our model consistently generates output that respects the provided sentence, further showing the capability of FashionGAN in retaining structural coherence.

Comparison with One-Step GAN Baselines. We provide a qualitative comparison with One-Step variants in Fig. 8. As shown in the figure, our approach achieves better visual quality with fewer artifacts and more consistent human shape.

Comparison with the Non-Compositional Baseline. We show the results in Fig. 9. Our approach provides clearer clothing regions while much fewer visual artifacts and noise, outperforming the baseline approach.

Comparison with the 2D Non-Parametric Baseline. We compare with this conventional baseline by retrieving an exemplar from a large database by text and perform Poisson image blending to apply the new outfit on the wearer. Results are shown in Fig. 10. Due to shape inconsistency between the exemplar and wearer’s body, the rendering results are not satisfactory.

4.3. User Study

Evaluating the Generated Segmentation Maps. A total of 50 volunteers participated in our user study. Our goal is to examine the quality of the intermediate segmentation maps generated by the first stage of FashionGAN. To



Figure 8. **Comparison with one-step GAN baselines.** (We use abbreviation here. One-8-7 refers to One-Step-8-7 and One-8-4 refers to One-Step-8-4). Each row represents one approach and each column shares the same text input. While these baseline approaches can also keep the human shape to some extent, the artifact presented in their results are significantly stronger than our results. The shape of the generated persons are also less consistent for baseline approaches.

this end, we provided the human segmentation map of the original photograph and the generated map, *i.e.*, a pair of maps for each test case, and asked participants to determine which map looked more realistic and genuine. A higher number of misclassified test cases implies a better quality of the generated maps. As only FashionGAN would produce such intermediate segmentation map, we thus only conduct this experiment with our approach. For the total of 8,979 test cases, the participants misclassified 3,753 of them (the misclassification rate is 42%). This is significant as our



Figure 9. **Comparison with the non-compositional baseline.** All the images are generated from the same human shape map from the previous stage. The upper row comes from the proposed compositional approach while the bottom row comes from the baseline method.



Figure 10. **Representative failure cases for the 2D non-parametric baseline.** Shape inconsistency between the retrieved exemplar and body causes mismatch between the new outfit and the wearer.

segmentation maps fooled most of the participants, whose ratings were close to random guessing.

Evaluating the Generated Photos. The same group of volunteers were asked to provide a ranking of the generated images produced by FashionGAN as well as the results from three baseline approaches, namely, ‘One-Step-8-7’, ‘One-Step-8-4’, and ‘Non-Compositional’. In addition, we also compared against the 2D non-parametric approach. During the user study, each participant was provided with the original image and the corresponding sentence description. The participants were asked to rank the quality of the generated images with respect to the relevance to the sentence description and the texture quality.

All the 8,979 test images were evenly and randomly assigned to these participants. We summarize various statistics in Table 2 and the frequency statistics for each rating in Fig. 11. For each approach, we computed the average ranking (where 1 is the best and 5 is the worst), standard deviation, and the frequency of being assigned with each ranking. We can observe that most of the high ranks go to our approach, which indicates that our solution achieves the best visual quality and relevance to the text input.

Table 2. A user study that evaluates the quality of generated images. A smaller number indicates a higher rank.

	Mean Ranking	Std Ranking
One-Step-8-7	4.027	0.894
One-Step-8-4	4.097	0.993
Non-Compositional	3.045	1.193
2D Non-Parametric	2.286	1.002
Ours	1.544	0.869

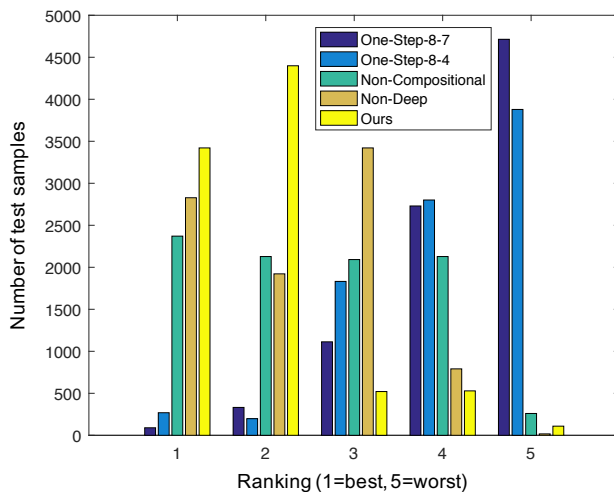


Figure 11. Detailed user study results. Each user rates 1 to 5 to each image they are assigned, and we show which methods each of these ratings go to. Rank 1 is the best and 5 is the worst.

5. Conclusion

We presented a novel approach for generating new clothing on a wearer based on textual descriptions. We designed two task-specific GANs, the shape and the image generators, and an effective spatial constraint in the shape generator. The generated images are shown to contain precise regions that are consistent with the description, while keeping the body shape and pose of a person unchanged. Quantitative and qualitative results outperform the baselines.

The results generated are limited by the current database we adopted. Our training set contains images mostly with a plain background as they were downloaded from on-line shopping sites (*i.e.*, <http://www.forever21.com/>). Hence the learning model is biased towards such a distribution. In fact we do not assume any constraints or post-processing of the background. We believe that our model can also render textured background if the training set contains more images with textured background. The background distribution will be captured by the latent vector z .

Acknowledgement: This work is supported by SenseTime Group Limited and the General Research Fund sponsored by the Research Grants Council of the Hong Kong SAR (CUHK 14224316, 14209217).

References

- [1] G. Gkioxari, R. Girshick, and J. Malik. Contextual action recognition with R*CNN. In *ICCV*, 2015.
- [2] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014.
- [3] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- [4] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pages 694–711. Springer, 2016.
- [5] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [6] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 2017.
- [7] X. Liang, C. Xu, X. Shen, J. Yang, S. Liu, J. Tang, L. Lin, and S. Yan. Human parsing with contextualized convolutional neural network. In *ICCV*, 2015.
- [8] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. DeepFashion: Powering robust clothes recognition and retrieval with rich annotations. In *CVPR*, 2016.
- [9] A. Nguyen, A. Dosovitskiy, J. Yosinski, T. Brox, and J. Clune. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. In *NIPS*, 2016.
- [10] D. Protopsaltou, C. Luible, M. Arevalo, and N. Magnenat-Thalmann. A body and garment creation method for an internet based virtual fitting room. In *Advances in Modelling, Animation and Rendering*, pages 105–122. Springer, 2002.
- [11] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [12] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text-to-image synthesis. In *ICMR*, 2016.
- [13] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. *TOG*, 23(3):309–314, 2004.
- [14] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training GANs. In *NIPS*, 2016.
- [15] X. Wang and A. Gupta. Generative image modeling using style and structure adversarial networks. In *ECCV*, 2016.
- [16] S. Yang, T. Ambert, Z. Pan, K. Wang, L. Yu, T. Berg, and M. C. Lin. Detailed garment recovery from a single-view image. *arXiv preprint arXiv:1608.01250*, 2016.
- [17] D. Yoo, N. Kim, S. Park, A. S. Paek, and I. S. Kweon. Pixel-level domain transfer. In *ECCV*, 2016.
- [18] H. Zhang, T. Xu, H. Li, S. Zhang, X. Huang, X. Wang, and D. Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, 2017.
- [19] S. Zhou, H. Fu, L. Liu, D. Cohen-Or, and X. Han. Parametric reshaping of human bodies in images. *TOG*, 29(4):126, 2010.