

Shape-Guided Clothing Warping for Virtual Try-On

Xiaoyu Han
xyhan@stu.hit.edu.cn
Harbin Institute of Technology
Weihai, China

Chenyang Wang
c.wang@stu.hit.edu.cn
Harbin Institute of Technology
Weihai, China

Shunyuan Zheng
sawyer0503@hit.edu.cn
Harbin Institute of Technology
Weihai, China

Xin Sun
sunxintyc@hit.edu.cn
Harbin Institute of Technology
Weihai, China

Zonglin Li
zonglin.li@hit.edu.cn
Harbin Institute of Technology
Weihai, China

Quanling Meng*
quanling.meng@hit.edu.cn
Harbin Institute of Technology
Weihai, China



Figure 1: Virtual try-on typically deforms clothing to fit the person's body, which is then combined with the person image to obtain try-on results. Compared with existing TPS-based methods (e.g., VITON-HD [5]) and appearance flow based methods (e.g., HR-VTON [28]), our method excels in capturing the clothing shape conforming to the person's body.

Abstract

Image-based virtual try-on aims to seamlessly fit in-shop clothing to a person image while maintaining pose consistency. Existing methods commonly employ the thin plate spline (TPS) transformation or appearance flow to deform in-shop clothing for aligning with the person's body. Despite their promising performance, these methods often lack precise control over fine details, leading to inconsistencies in shape between clothing and the person's body as well as distortions in exposed limb regions. To tackle these challenges, we propose a novel shape-guided clothing warping method for virtual try-on, dubbed SCW-VTON, which incorporates global shape constraints and additional limb textures to enhance the realism and consistency of the warped clothing and try-on results. To integrate global shape constraints for clothing warping, we devise a dual-path clothing warping module comprising a shape path and a flow path. The former path captures the clothing shape aligned with the person's body, while the latter path leverages the mapping

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MM '24, October 28-November 1, 2024, Melbourne, VIC, Australia
© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0686-8/24/10
<https://doi.org/10.1145/3664647.3680756>

between the pre- and post-deformation of the clothing shape to guide the estimation of appearance flow. Furthermore, to alleviate distortions in limb regions of try-on results, we integrate detailed limb guidance by developing a limb reconstruction network based on masked image modeling. Through the utilization of SCW-VTON, we are able to generate try-on results with enhanced clothing shape consistency and precise control over details. Extensive experiments demonstrate the superiority of our approach over state-of-the-art methods both qualitatively and quantitatively.

CCS Concepts

• Computing methodologies → Reconstruction.

Keywords

Virtual Try-on; Image Synthesis; Appearance Flow

ACM Reference Format:

Xiaoyu Han, Shunyuan Zheng, Zonglin Li, Chenyang Wang, Xin Sun, and Quanling Meng. 2024. Shape-Guided Clothing Warping for Virtual Try-On. In *Proceedings of the 32nd ACM International Conference on Multimedia (MM '24)*, October 28-November 1, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3664647.3680756>

1 Introduction

In recent years, the e-commerce industry has experienced rapid development, leading to an increasing number of consumers purchasing clothing online. To provide online customers with a shopping experience that rivals in-store try-on, significant attention has been

devoted to exploring the virtual try-on technology. Virtual try-on can be broadly categorized into 3D-based [12, 34, 36, 42, 43, 47] and image-based methods [7–10, 14, 15, 18, 22, 24, 30, 45, 48–51, 54, 55, 57, 59, 60], with the latter garnering more interest due to its lightweight data and wider applicability. This study will specifically delve into image-based virtual try-on.

Among the key processing steps in image-based virtual try-on, clothing warping stands out as particularly challenging. Its objective is to deform the target clothing to align with the person's body, which is vital for maintaining consistency between the clothing shape and the person's posture. Early methods [15, 33, 45, 55] adopt the thin plate spline (TPS) transformation [3] to achieve this deformation. However, these TPS-based methods, constrained by limited degrees of freedom, prove inadequate when significant geometric deformations are needed [6, 14], such as the misaligned clothing generated by VITON-HD [5] in Figure 1. To overcome this limitation, recent studies [1, 6, 14, 31, 53] introduce appearance flow [61] as a solution, which provides dense pixel-level predictions for deforming from the source clothing to the target one, thereby enhancing the accuracy of the resulting warped clothing. While these methods have demonstrated promising performance in virtual try-on, a notable limitation remains in their current inability to exert precise control over the details of the obtained warped clothing. Firstly, the lack of the appropriate constraint on shape during pixel-level prediction may lead to inconsistencies in the shape of the warped clothing compared to the person's body. For instance, as shown in the HR-VTON [28] column of Figure 1, the generated clothing edges appear torn and inconsistent. Secondly, distortions may occur partially in the try-on result, especially in the limb regions not covered by clothing. This is attributed to the absence of detailed information guidance for these parts, as existing methods often use the person image with masked clothing and limb regions as the input for try-on synthesis.

In this study, we direct our attention towards addressing these underexplored challenges. We present SCW-VTON, a shape-guided clothing warping method for virtual try-on, which incorporates global shape constraints and additional limb texture references to enhance the realism of generated warped clothing and try-on results. Initially, we leverage global shape constraints to facilitate the clothing warping phase, alleviating discrepancies between the warped clothing shape and the person's body. We design a dual-path clothing warping module comprising a shape path and a flow path. The shape path first predicts the shape of the target clothing, facilitating the subsequent exploration of the mapping between the pre- and post-deformation clothing shapes using a set of shape-guided cross-attention blocks. The acquired mapping is then integrated into a flow path as global shape constraints to steer the estimation of appearance flow. This flow is subsequently applied to in-shop clothing to create the target warped clothing. It is worth mentioning that introducing additional global shape constraints frees our method from reliance on the original input shape, which implies that our method is also capable of transferring textures from a shapeless texture map to the person image while preserving their distribution. Furthermore, we employ a co-training strategy for the two paths during training, which positively affects the weight update of the flow estimation model and further enhances the stability

and accuracy of appearance flow. In the subsequent process of try-on synthesis, to alleviate distortions in limb regions, we develop an additional limb reconstruction network based on the masked image modeling method. Specifically, we first derive a limb map from the source person image. Then, we employ an autoencoder to take the masked partial limb component as input to learn latent limb representations, from which we reconstruct realistic limb textures at specific locations. Through the combination of shape-guided appearance flow and reconstructed limb textures, our approach can ultimately produce try-on results with enhanced clothing shape consistency and precise control over details.

The contributions of this paper are summarized as follows:

- We introduce a shape-guided clothing warping method for virtual try-on that incorporates global shape constraints on clothing deformation, resulting in realistic warped clothing that conforms accurately to the person's body.
- We design a limb reconstruction network to provide precise guidance on generating limb regions of the try-on result, effectively addressing issues such as performance degradation and distortions during the try-on synthesis process.
- Extensive experiments demonstrate the superior performance of our proposed SCW-VTON compared to existing state-of-the-art methods for virtual try-on.

2 Related Work

2.1 Appearance Flow

Appearance flow is applied to predict a 2D vector field and warp the source image to the target based on the similarity in appearance, which is proposed by Zhou et al. [61] to solve the problem of novel view synthesis. In recent years, appearance flow has been widely applied in various fields of computer vision. In image inpainting tasks [39, 52], appearance flow is used to propagate pixels from source to missing regions, enhancing realism in generated contents. Additionally, appearance flow also enables human pose transfer [4, 29, 38, 58], synthesizing novel poses by warping the feature representations of the human body. Besides, virtual try-on has attracted widespread attention in recent years and many methods [1, 10, 14, 16] introduce the appearance flow to deform the in-shop clothing to achieve the alignment with the person's body.

2.2 Image-based Virtual Try-on

Clothing Warping. Clothing Warping is typically a key processing step of most image-based virtual try-on methods. Early methods [15, 33, 45, 55] primarily achieve this through thin-plate spline (TPS) transformation. However, TPS-based methods are unable to accurately handle large geometric deformation due to the limited degrees of freedom. Therefore, ClothFlow [14] proposes an appearance flow based method to perform clothing warping with pixel-level displacements. Then, Zflow [6] adopts gated appearance flow to further stabilize the deformation. SAL-VTON [53] links the clothing with the person's body with semantically associated landmarks and proposes the local flow to alleviate the misalignment. However, considering the potential drastic deformation caused by appearance flow, we introduce extra global shape constraints based on a cross-attention mechanism, which warrants a more stable flow and makes the warped clothing match a consistent shape.

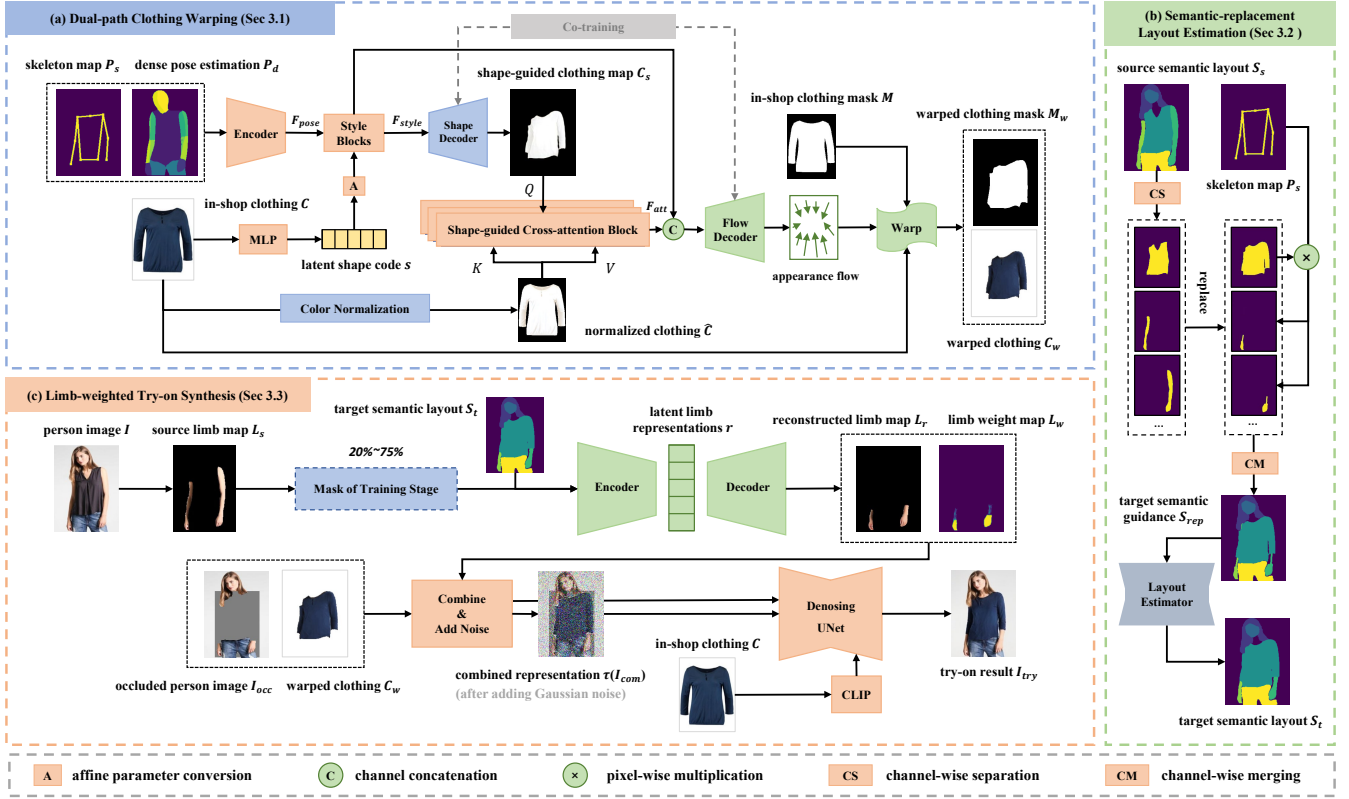


Figure 2: The overview of SCW-VTON. (a) Warped clothing C_w is produced through a dual-path architecture, where the cross-attention features F_{att} are created as global shape constraints to assist in estimating appearance flow. (b) Based on a semantic-replacement strategy, a target semantic layout S_t that describes the person wearing target clothing is predicted. (c) A pre-trained autoencoder produces the reconstructed limb map L_r from latent limb representations r , which is then combined with the occluded person image I_{occ} and warped clothing C_w to produce the final try-on result I_{try} with a diffusion model.

Try-on Synthesis. Another important step is try-on synthesis, which refers to synthesizing the target clothing and the person image to produce the try-on result. Benefiting from the powerful generation capability of the diffusion model [20, 40], some existing methods [2, 11, 26, 35, 62] introduce it as the final component of try-on synthesis. DCI-VTON [11] treats virtual try-on as an inpainting task and incorporates the diffusion model to refine a coarse result. LaDI-VTON [35] employs textual inversion technique in virtual try-on and proposes EMASC modules to enhance synthesis quality. StableVITON [26] proposes a zero cross-attention block used in the pre-trained diffusion decoder to learn the semantic correspondence between the clothing and the person for achieving better detail preservation. However, most methods require the synthesis network to reconstruct limb textures in the absence of corresponding reference cues from the person image, leading to the degradation of performance and distortions in try-on results. We attempt to alleviate this problem by creating an additional limb reconstruction network based on masked image modeling.

3 Proposed Method

As shown in Figure 2, SCW-VTON consists of three modules: 1) a dual-path clothing warping module, 2) a semantic-replacement

layout estimation module, and 3) a limb-weighted try-on synthesis module. The dual-path clothing warping module produces the warped clothing C_w and the warped clothing mask M_w while the semantic-replacement layout estimation module generates the target semantic layout S_t . Based on these acquired results, the limb-weighted try-on synthesis module ultimately generates the try-on result I_{try} through a diffusion model.

3.1 Dual-path Clothing Warping

Figure 2 (a) shows the schematic of the dual-path clothing warping module, which can be divided into a shape path, shape-guided cross-attention blocks, and a flow path. We begin by introducing the shape characteristics, which will be used in the shape path and shape-guided cross-attention blocks.

Shape Characteristics. The first consideration is how to obtain the shape characteristics of clothing pre- and post-deformation, as this is essential for constructing the mapping between them. In this section, we design a color normalization strategy to obtain the shape characteristics of clothing before deformation firstly, which explicitly normalizes each color channel of in-shop clothing C to discard its raw color while keeping the overall gradient difference. Specifically, as shown in Figure 3, we first calculate the average




Figure 3: Shape-guided cross-attention blocks of SCW-VTON.

color value ξ^k for each channel k of in-shop clothing C . Then we subtract ξ^k from each pixel value within the clothing regions of C . Finally, we add in-shop clothing mask M to the subtracted result and get the normalized clothing \hat{C} , which can reflect the shape characteristics of in-shop clothing more intuitively while excluding interfering information. The above process can be expressed as:

$$\xi^k = \frac{\sum_{i=0}^{H_C} \sum_{j=0}^{W_C} (C^{k,i,j} \odot M^{i,j})}{\sum_{i=0}^{H_C} \sum_{j=0}^{W_C} M^{i,j}}, \quad (1)$$

$$\hat{C}^k = \mathcal{T}(C^k - \xi^k + M), \quad (2)$$


where \odot is the element-wise multiplication, i and j is the position of a sample pixel, H_C and W_C are the height and width of C , respectively. $\mathcal{T}(\cdot)$ is a truncated function ensuring the output value falls within the range of zero and one.

Shape Path. The subsequent goal is to obtain the shape characteristics of clothing after deformation, which are then used to match the normalized clothing \hat{C} . We first adopt a CNN-based encoder to extract human pose features F_{pose} from the skeleton map P_s and dense pose estimation P_d , where P_s is obtained by connecting 9 key points of the person's body and P_d is from [13]. Although F_{pose} already contains basic information that matches the person's body, it is also essential to ensure that the design and style of the generated clothing are consistent with the in-shop clothing C . Inspired by [23], we adopt a multilayer perceptron (MLP) to embed in-shop clothing C into the latent space as a global style code s , which is used to calculate a set of affine transformation parameters in style blocks to adjust the pose features F_{pose} and get the stylized shared features F_{style} for the shape decoder and the subsequent flow decoder. This process can be formulated as:

$$F_{style}^{k,i,j} = \gamma^{k,i,j}(s) \frac{F_{pose}^{k,i,j} - \mu^k}{\sigma^k} + \delta^{k,i,j}(s), \quad (3)$$

where $F_{style}^{k,i,j}$ is a particular sample of F_{style} at location (k, i, j) , $\gamma(\cdot)$ and $\delta(\cdot)$ are the convolution operations that convert the input into affine parameters, μ^k and σ^k are the mean and standard deviation of pose features F_{pose}^k , respectively. μ^k is calculated as:

$$\mu^k = \frac{1}{H_F W_F} \sum_{i=0}^{H_F} \sum_{j=0}^{W_F} F_{pose}^{k,i,j}, \quad (4)$$

Figure 4: Compared to the binary mask M_w , our shape-guided clothing map C_s contains detailed shape characteristics of clothing such as the basic appearance, shadows, and wrinkles.

where H_F and W_F are the height and width of F_{pose} , respectively. σ^k is calculated as:

$$\sigma^k = \sqrt{\frac{1}{H_F W_F} \sum_{i=0}^{H_F} \sum_{j=0}^{W_F} (F_{pose}^{k,i,j} - \mu^k)^2}. \quad (5)$$

Then, we adopt a shape decoder to up-sample the shared features F_{style} and generate a shape-guided clothing map C_s , where we employ our color normalization strategy again to get the ground truth of C_s in the training phase (please refer to Section 3.4 for more details). Note that C_s is not the binary mask of the warped clothing, as shown in Figure 4, compared to the binary mask M_w obtained by directly applying flow to the in-shop clothing mask M , C_s encompasses detailed shape characteristics of clothing after deformation such as the basic appearance, shadows, and wrinkles. **Shape-guided Cross-attention Blocks.** After obtaining the normalized clothing \hat{C} and the shape-guided clothing map C_s that respectively reflect the shape characteristics of clothing before and after deformation, we adopt cross-attention blocks to explore the mapping between them, which is subsequently used to constrain the estimation of appearance flow. As shown in Figure 3, C_s and \hat{C} are fed into three separate encoders to get Q, K, V , respectively:

$$Q = Fla(E_q(C_s)), K = Fla(E_k(\hat{C})), V = Fla(E_v(\hat{C})), \quad (6)$$

where E_q, E_k, E_v are encoders to get the corresponding components of the cross-attention mechanism, and $Fla(\cdot)$ is the flattening operation for extracted feature maps. Next, for the i -th cross-attention block, its output can be represented as:

$$f_{att}^i = softmax\left(\frac{QA_q^i(KA_k^i)^T}{\sqrt{d}}\right)VA_v^i, \quad (7)$$

where A_q, A_k, A_v represent linear layers, and d is the dimension of KA_k^i . Finally, all the outputs of cross-attention blocks are integrated to form the cross-attention features F_{att} .

Flow Path. The essence of appearance flow is the pixel-wise displacement, which does not involve the degradation of image quality caused by down-sampling, so a warped result applied by appearance flow has the ability to retain detailed textures. However, directly estimating appearance flow without global shape constraints may lead to unnatural drastic deformations. Therefore, instead of directly decoding the shared features F_{style} to get the appearance flow, we introduce F_{att} as extra global shape constraints to perform the shape-guided warping. Specifically, the inputs to the flow decoder consist of two parts: one is the shared features F_{style} , and the other is the cross-attention features F_{att} . They are concatenated




Figure 5: Comparison of the try-on result w/o and w/ the reconstructed limb map L_r and the limb weight map L_w .

and fed into the flow encoder together to perform pixel-wise regression and estimate the appearance flow. Note that at the training stage, the weight parameters of the shape decoder and the flow decoder are updated simultaneously. This enables the flow decoder to collaborate with the shape decoder seamlessly and implicitly, thereby predicting the appearance flow with a more consistent shape. Finally, the appearance flow is applied to the in-shop clothing C and the in-shop clothing mask M , getting the warped clothing C_w and the corresponding warped clothing mask M_w .

3.2 Semantic-replacement Layout Estimation

Given the warped clothing mask M_w , the skeleton map P_s , and the source semantic layout S_s , the goal of the semantic-replacement layout estimation module is to produce the target semantic layout S_t , which describes the person wearing new clothing. S_t is utilized to provide the location information for limb reconstruction in the following process of try-on synthesis.

Since it is almost unavailable to acquire the training data about a person wearing two different clothing in a fixed pose, it is a necessary and common practice to acquire person representations that discard the person's original clothing information to train the network in a self-supervised way. We use a semantic-replacement strategy to perform the above discarding during predicting S_t . As shown in Figure 2 (b), the source semantic layout S_s is separated into a multi-channel binary parsing map and each channel corresponds to clothing or a part of the person's body. We utilize the warped clothing mask M_w and a portion of the skeleton map P_s to replace the original clothing and limb channels to get the target semantic guidance S_{rep} , which ensures semantic continuity and pose consistency while preventing negative interference with the network training caused by the person's original clothing. Subsequently, we use a UNet [41] model as the semantic layout estimator, which takes S_{rep} as the input and predicts the target semantic layout S_t . More detailed descriptions of this module can be found in the supplementary material.

3.3 Limb-weighted Try-on Synthesis

After obtaining the warped clothing C_w and the target semantic layout S_t , our final goal is to generate the try-on result. However, during this process, reconstructing limb textures is a challenging task, as the lack of sufficient reference information may lead to inconsistency between the distribution of the reconstructed result and the original image.

Limb Reconstruction. Therefore, we propose a separate limb reconstruction network, which focuses on learning compact latent representations of the person's limbs through masked image

modeling, aiming to enhance the realism and consistency of the reconstructed limb textures with the original data. As depicted in Figure 2 (c), we first extract the source limb map L_s from the person image I . Inspired by [17], we randomly mask L_s by 20%~75% and obtain the masked limb map before an autoencoder. On the one hand, data pairs with the asymmetric input and output are created in the training stage, enhancing the ability of the autoencoder to reconstruct limb textures. On the other hand, random masking encourages the network to learn more compact limb representations. Note that the mask operation is only set in the training phase, while during testing, the autoencoder directly takes L_s as input. Besides, the target semantic layout S_t is another input of the autoencoder, which is used to provide location information for limb reconstruction. The autoencoder is pre-trained individually and is capable of outputting the reconstructed limb map L_r and the limb weight map L_w , where the latter is a correlation weight to adaptively modify the limb regions that may be prone to performance degradation and distortions (as shown in Figure 5).

Try-on Synthesis. Based on L_w , L_r is then combined with the occluded person image I_{occ} (obtained by occluding the clothing and limb regions of I) and the warped clothing C_w to get the combined representation I_{com} . Inspired by [11], we adopt the pre-trained diffusion model as the final try-on synthesis network, utilizing the combined representation I_{com} as the input to produce the try-on result I_{try} . Besides, a pre-trained CLIP [37] image encoder is employed to extract additional conditions, which are injected into the Denoising UNet through the cross-attention mechanism to guide the generation of the diffusion model.

3.4 Training Objectives

Dual-path Clothing Warping. To train the dual-path clothing warping module, we first adopt the reconstruction loss to constrain the pixel-wise value of C_s and C_w , which is formulated as:

$$l_{rec} = \|C_s - \hat{C}_{gt}\|_1 + \|C_w - C_{gt}\|_1, \quad (8)$$

where C_{gt} is the ground truth of the warped clothing C_w , it is extracted from the person image I . \hat{C}_{gt} is obtained by applying the color normalization strategy to C_{gt} . Besides, we adopt the perceptual loss proposed in [25] to calculate the distance of the features extracted by the VGG-19 [44] network:

$$l_{per} = \sum_{i=1}^5 (\|\phi_i(C_s) - \phi_i(\hat{C}_{gt})\|_1 + \|\phi_i(C_w) - \phi_i(C_{gt})\|_1), \quad (9)$$

where $\phi_i(\cdot)$ denotes the feature maps of the i -th layer in the pre-trained perception network. Furthermore, we adopt a mask loss to constrain the warped clothing mask M_w , which is formulated as:

$$l_{mask} = \|M_w - M_{gt}\|_1, \quad (10)$$

where M_{gt} is the mask of C_{gt} . Overall, the loss of the dual-path clothing warping module is represented as:

$$l_{warp} = l_{rec} + \lambda_{per} l_{per} + l_{mask}, \quad (11)$$

where λ_{per} is used to balance the weights of these losses.

Semantic-replacement Layout Estimation. We use a weighted cross-entropy loss to supervise the training process of the semantic



Figure 6: Qualitative results of baseline methods and our SCW-VTON on the VITON [15] dataset.

layout estimator, which is expressed as:

$$l_{sem} = -\frac{1}{n} \sum_{i=0}^n \sum_{j=0}^c w_j S_s^{i,j} \log(S_t^{i,j}), \quad (12)$$

where n denotes the number of samples, c is the number of channels of S_s and S_t , and w_j is the loss weight in the j -th class channel.

Limb-weighted Try-on Synthesis. Finally, following [11], the loss of the limb-weighted try-on synthesis l_{syn} is designed in two parts, which can be represented as:

$$l_{syn} = \lambda_{vgg} l_{vgg} + l_{ldm}. \quad (13)$$

l_{vgg} is similar to l_{per} :

$$l_{vgg} = \sum_{i=1}^5 \|\phi_i(I_{try}) - \phi_i(I)\|_1, \quad (14)$$

and l_{ldm} is formulated as:

$$l_{ldm} = \|\epsilon - \epsilon_\theta(\tau(\mathcal{E}(I_{com})), \mathcal{E}(I_{com}), m, CLIP(C), t)\|_2^2, \quad (15)$$

where \mathcal{E} is a pre-trained encoder belonging to the diffusion model, which embeds the images from image space to latent space. $\tau(\cdot)$ is the operation of adding noise, m is the mask used to occlude the person image I , and t is the timestamp.

4 Experiment

4.1 Datasets

We conduct main experiments on the public virtual try-on benchmark dataset VITON [15], which contains 14,221 data pairs for training and 2,032 data pairs for testing. Also, we carry out the comparative experiments under the higher resolution on the VITON-HD [5] dataset, which contains 11,647 data pairs for training and 2,032 data pairs for testing.

4.2 Implementation Details

We train three modules of SCW-VTON independently. For the dual-path clothing warping module and the semantic layout estimation module, they are both trained for 30 epochs and optimized by Adam [27] with $\beta_1=0.5$ and $\beta_2=0.999$. The learning rate is fixed at 0.0001 in the first half of training and then linearly decays to zero in the remaining steps. We set the hyper-parameters as $\lambda_{per} = 5$, $w_0 = w_1 = w_2 = w_6 = 1$, and $w_3 = w_4 = w_5 = 3$. For the limb-weighted try-on synthesis module, following [11], we use AdamW [32] optimizer with the learning rate of 0.00001 to train this module for 50 epochs, and the hyper-parameters λ_{vgg} is set to 0.0001.

4.3 Qualitative Results

We first compare our SCW-VTON with existing virtual try-on methods ACGPN [55], SDAFN [1], RMGN [31], and DOC-VTON [56] on the VITON [15] dataset qualitatively. In Figure 6, we divide the comparison into two groups, where the left group focuses on the clothing and the right group focuses on the person's limbs. On the one hand, most baseline methods struggle to handle clothing with dense and complex textures. Take the first row of the left group in Figure 6 as an example, the logo after deformation presents a large range of distortions in the results of RMGN [31] and DOC-VTON [56], while the results obtained by ACGPN [55] and SDAFN [1] appear blurry. In comparison, benefiting from the proposed shape constraints on clothing warping, the global shape and pose consistency of clothing is warranted in our results, leading to remarkable realism. On the other hand, it is also challenging to fit in-shop clothing into a person image with a complex pose, which mainly involves limb occlusion and rotation. The right group of Figure 6 showcases the inferior performance of baseline methods in these difficult cases. For instance, in the first row of the right group, ACGPN [55] and SDAFN [1] generate distorted and unnatural limbs, while RMGN [31] and DOC-VTON [56] disrupt the continuity of limb textures. Compared to baseline methods, we achieve more reasonable effects, which are attributed to the consistent clothing shape and realistic limb textures acquired by our method. Qualitative experiments on the VITON-HD [5] dataset are conducted in Figure 7, where the results of HR-VTON [28], SAL-VTON [53], DCI-VTON [11], StableVITON [26], and our SCW-VTON are represented. Similarly, most baseline methods struggle to align clothing with the person's body while maintaining the integrity of the clothing design. For example, in the first row on the left side of Figure 7, most baseline methods fail to accurately estimate the correct position of the waistband, and even lose the texture information of the waistband after the clothing deformation. In contrast, our method can preserve the waistband intact and deform it to the appropriate position according to the person's posture. These results validate the robust performance of our method even with an increased resolution. Additional qualitative experimental results are provided in the supplementary material.

4.4 Quantitative Results

We conduct quantitative experiments both in paired data setting and unpaired data setting, denoting that a person wears original clothing and new clothing after try-on, respectively. We use Structural Similarity (SSIM) [46] and Peak Signal to Noise Ratio (PSNR) [21] as



Figure 7: Qualitative results generated by baseline methods and our SCW-VTON on the VITON-HD [5] dataset.

Table 1: Quantitative comparisons on the VITON [15] and VITON-HD [5] datasets. The percentage results in the last column are displayed as a/b, with a and b representing the preference for the baseline method and our method, respectively.

Method	Dataset	Warped Clothing		Try-on Results			User Study
		SSIM (\uparrow)	PSNR (\uparrow)	FID (\downarrow)	SSIM (\uparrow)	PSNR (\uparrow)	
ACGPN [55]	VITON	0.8711	22.38	12.77	0.8454	23.12	12.83% / 87.17%
PL-VTON [16]		0.8434	19.94	11.68	0.8507	24.73	18.67% / 81.33%
SDAFN [1]		-	-	10.83	0.8399	23.49	31.33% / 68.67%
RMGN [31]		0.8588	21.28	10.52	0.8633	24.95	29.83% / 70.17%
DOC-VTON [56]		0.8531	19.05	9.43	0.8323	22.07	33.02% / 66.98%
SCW-VTON (ours)		0.8839	23.31	8.89	0.8897	26.78	references
VITON-HD [5]	VITON-HD	0.8641	18.29	14.05	0.8497	21.19	23.08% / 76.92%
HR-VTON [28]		0.8622	18.41	11.69	0.8657	22.38	29.28% / 70.72%
SAL-VTON [53]		-	-	9.64	0.8798	23.29	36.72% / 63.28%
GP-VTON [49]		0.8905	21.76	9.52	0.8735	23.31	36.08% / 63.92%
LaDI-VTON [35]		0.8897	21.88	9.57	0.8638	22.52	31.80% / 68.20%
DCI-VTON [11]		0.8841	21.15	9.67	0.8749	23.07	35.20% / 64.80%
StableVTON [26]		-	-	9.45	0.8678	23.48	39.32% / 60.68%
SCW-VTON (ours)		0.8971	22.47	8.96	0.8829	23.98	references

metrics for the paired data setting and Fréchet Inception Distance (FID) [19] for the unpaired data setting. To perform a more comprehensive evaluation, we compute metrics both for the process of clothing warping and overall try-on. Table 1 summarizes the quantitative results of our method and baseline methods on two datasets, which indicates that our SCW-VTON outperforms all the baseline methods in SSIM, PSNR, and FID, both in terms of the warped clothing and try-on result.

4.5 User Study

We conduct a user study with 30 recruited volunteers to evaluate the visual effect of our method, where SCW-VTON is compared with each baseline method in an A/B manner. Specifically, 200 in-shop clothing images and 200 person images are randomly selected from the testing set in each comparison, which are used to output the try-on results by SCW-VTON and another baseline method to form 200 pairwise result groups. The volunteers are asked to consider the rationality of two try-on results from each group carefully and choose the better one. The results are summarized in the last column Table 1, it is evident that our SCW-VTON always provides a better visual experience in each A/B comparison.

4.6 Ablation Studies

Global Shape Constraints. To assess the impact of the cross-attention features F_{att} as global shape constraints, we design a variant SCW-VTON* that removes shape-guided cross-attention blocks, where the flow decoder only takes the shared features F_{style} as the input. We first compare the variant and our full model in Figure 8 (a). It is evident that SCW-VTON* is difficult to align well with the person’s body, especially at the edge of clothing, while SCW-VTON estimates the clothing shape more accurately, and effectively addresses the problem of drastic deformation occurring in the variant. Quantitative results of SCW-VTON and SCW-VTON* are reported in Table 2, which shows that SCW-VTON outperforms the variant across all metrics. The supplementary material provides more analysis and results on the global shape constraints.

Co-training. We design a variant SCW-VTON* to verify the effect of the co-training in the dual-path clothing warping module. For this variant, we only remove the co-training strategy implemented in the shape decoder and the flow decoder, while keeping all other components consistent with the full model. SCW-VTON and SCW-VTON* are quantitatively compared in Table 2, while Figure 8 (a) illustrates their differences in visual effects. It can be seen that

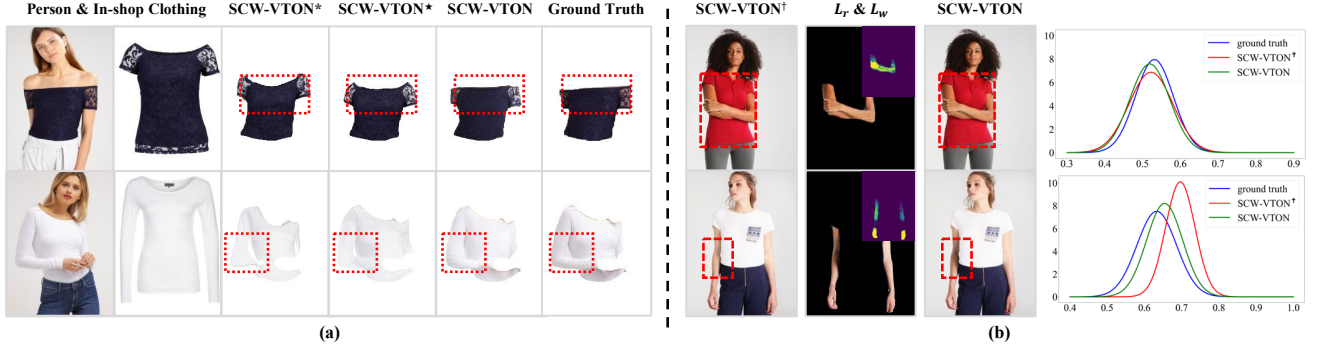


Figure 8: Qualitative results of ablation studies. (a) The variant SCW-VTON* removes shape-guided cross-attention blocks, while SCW-VTON* only removes the co-training strategy. (b) SCW-VTON† produces try-on results without the limb reconstruction network. The curve graphs on the right represent the pixel value distribution of corresponding samples in the limb regions.

Table 2: Quantitative results of ablation studies on the VITON [15] dataset.

Method	Config	Warped Clothing		Try-on Results		
		SSIM (↑)	PSNR (↑)	FID (↓)	SSIM (↑)	PSNR (↑)
SCW-VTON*	w/o global shape constraints	0.8423	19.95	10.16	0.8639	25.35
SCW-VTON*	w/o co-training strategy	0.8792	22.18	9.44	0.8723	26.31
SCW-VTON†	w/o limb reconstruction network	-	-	9.24	0.8740	25.75
SCW-VTON (ours)	full model	0.8839	23.31	8.89	0.8897	26.78



Figure 9: Examples of try-on results by our SCW-VTON based on the pure texture map input.

although some misalignment issues in SCW-VTON* are alleviated in SCW-VTON*, SCW-VTON* lacks realism in detailed regions compared to SCW-VTON. The supplementary material provides more analysis and results on the co-training strategy.

Limb Reconstruction Network. Finally, we propose a variant SCW-VTON† to evaluate the role of the limb reconstruction network in the process of try-on synthesis. We ablate the autoencoder that outputs the reconstructed limb map L_r , so the final try-on result is generated only based on the combined representation I_{com} without additional limb texture references. We compute the quantitative metrics for SCW-VTON†, which is also shown in Table 2. In addition, the qualitative comparison between this variant and

SCW-VTON is presented in Figure 8 (b), which illustrates that SCW-VTON can effectively solve the problems of performance degradation and distortions in SCW-VTON†, and make the distribution of reconstructed limb textures more consistent with the ground truth.

4.7 Potential Application

Besides the classic virtual try-on, we also try to explore other applications of our method. By introducing extra global shape constraints, our method is liberated from the restriction of generating try-on results solely based on specific input clothing shapes. For example, as shown in Figure 9, we can seamlessly transfer textures from a shapeless texture map to the person's body while preserving the original distribution. This versatility highlights that our method extends beyond conventional virtual try-on applications, potentially sparking innovative ideas for novel computer vision tasks in the fashion and clothing domain.

5 Conclusion

In this paper, we propose SCW-VTON, a novel shape-guided clothing warping method for virtual try-on. Based on a dual-path clothing warping module, our method incorporates additional global shape constraints on clothing deformation, resulting in realistic warped clothing that conforms accurately to the person's body. Besides, we design a limb reconstruction network based on masked image modeling to learn the compact latent limb representations and generate additional limb textures to refine the details in these regions through adaptive weighting. With the implementation of SCW-VTON, we produce try-on results with improved clothing shape consistency and precise control over details. Extensive experiments are conducted to demonstrate the superiority of our SCW-VTON over existing state-of-the-art methods.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 62072141).

References

- [1] Shuai Bai, Huiling Zhou, Zhikang Li, Chang Zhou, and Hongxia Yang. 2022. Single stage virtual try-on via deformable attention flows. In *Proceedings of the European Conference on Computer Vision*. Springer, 409–425.
- [2] Alberto Baldrati, Davide Morelli, Giuseppe Cartella, Marcella Cornia, Marco Bertini, and Rita Cucchiara. 2023. Multimodal garment designer: Human-centric latent diffusion models for fashion image editing. In *Proceedings of the IEEE International Conference on Computer Vision*. 23393–23402.
- [3] Fred L. Bookstein. 1989. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11, 6 (1989), 567–585.
- [4] Baoyu Chen, Yi Zhang, Hongchen Tan, Baocai Yin, and Xiuping Liu. 2021. PMAN: Progressive multi-attention network for human pose transfer. *IEEE Transactions on Circuits and Systems for Video Technology* 32, 1 (2021), 302–314.
- [5] Seunghwan Choi, Sunghyun Park, Minsoo Lee, and Jaegul Choo. 2021. Viton-hd: High-resolution virtual try-on via misalignment-aware normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 14131–14140.
- [6] Ayush Chopra, Rishabh Jain, Mayur Hemani, and Balaji Krishnamurthy. 2021. ZFlow: Gated Appearance Flow-Based Virtual Try-On with 3D Priors. In *Proceedings of the IEEE International Conference on Computer Vision*. 5433–5442.
- [7] Ruili Feng, Cheng Ma, Chengji Shen, Xin Gao, Zhenjiang Liu, Xiaobo Li, Kairi Ou, Deli Zhao, and Zheng-Jun Zha. 2022. Weakly supervised high-fidelity clothing model generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3440–3449.
- [8] Xin Gao, Zhenjiang Liu, Zunlei Feng, Chengji Shen, Kairi Ou, Haihong Tang, and Mingli Song. 2021. Shape controllable virtual try-on for underwear models. In *Proceedings of the 29th ACM International Conference on Multimedia*. 563–572.
- [9] Chongjian Ge, Yibing Song, Yuying Ge, Han Yang, Wei Liu, and Ping Luo. 2021. Disentangled cycle consistency for highly-realistic virtual try-on. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 16928–16937.
- [10] Yuying Ge, Yibing Song, Ruimao Zhang, Chongjian Ge, Wei Liu, and Ping Luo. 2021. Parser-Free Virtual Try-On via Distilling Appearance Flows. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8485–8493.
- [11] Junhong Gou, Siyu Sun, Jianfu Zhang, Jianlou Si, Chen Qian, and Liqing Zhang. 2023. Taming the Power of Diffusion Models for High-Quality Virtual Try-On with Appearance Flow. In *Proceedings of the 31st ACM International Conference on Multimedia*. 7599–7607.
- [12] Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. 2012. DRAPE: DRessing Any PErson. *ACM Transactions on Graphics* 31, 4 (2012), 1–10.
- [13] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. 2018. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7297–7306.
- [14] Xintong Han, Xiaojun Hu, Weilin Huang, and Matthew R Scott. 2019. ClothFlow: A Flow-Based Model for Clothed Person Generation. In *Proceedings of the IEEE International Conference on Computer Vision*. 10471–10480.
- [15] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S Davis. 2018. VITON: An Image-Based Virtual Try-On Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7543–7552.
- [16] Xiaoyu Han, Shengping Zhang, Qinglin Liu, Zonglin Li, and Chenyang Wang. 2022. Progressive Limb-Aware Virtual Try-On. In *Proceedings of the 30th ACM International Conference on Multimedia*.
- [17] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 16000–16009.
- [18] Sen He, Yi-Zhe Song, and Tao Xiang. 2022. Style-based global appearance flow for virtual try-on. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3470–3479.
- [19] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *Advances in Neural Information Processing Systems* 30 (2017), 6626–6637.
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* 33 (2020), 6840–6851.
- [21] Alain Hore and Djemel Ziou. 2010. Image Quality Metrics: PSNR vs. SSIM. In *International Conference on Pattern Recognition*. 2366–2369.
- [22] Chia-Wei Hsieh, Chieh-Yun Chen, Chien-Lung Chou, Hong-Han Shuai, Jiaying Liu, and Wen-Huang Cheng. 2019. FashionOn: Semantic-guided image-based virtual try-on with detailed human and clothing information. In *Proceedings of the 27th ACM International Conference on Multimedia*. 275–283.
- [23] Xun Huang and Serge Belongie. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*. 1501–1510.
- [24] Thibaut Issenbuth, Jérémie Mary, and Clément Calauzenes. 2020. Do not mask what you do not need to mask: a parser-free virtual try-on. In *Proceedings of the European Conference on Computer Vision*. Springer, 619–635.
- [25] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *Proceedings of the European Conference on Computer Vision*. Springer, 694–711.
- [26] Jeongho Kim, Gyojung Gu, Minho Park, Sunghyun Park, and Jaegul Choo. 2024. StableVITON: Learning Semantic Correspondence with Latent Diffusion Model for Virtual Try-On. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [27] Diederik P Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations* (2015).
- [28] Sangyun Lee, Gyojung Gu, Sunghyun Park, Seunghwan Choi, and Jaegul Choo. 2022. High-resolution virtual try-on with misalignment and occlusion-handled conditions. In *Proceedings of the European Conference on Computer Vision*. Springer, 204–219.
- [29] Yining Li, Chen Huang, and Chen Change Loy. 2019. Dense Intrinsic Appearance Flow for Human Pose Transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3693–3702.
- [30] Zhi Li, Pengfei Wei, Xiang Yin, Zejun Ma, and Alex C Kot. 2023. Virtual Try-On with Pose-Garment Keypoints Guided Inpainting. In *Proceedings of the IEEE International Conference on Computer Vision*. 22788–22797.
- [31] Chao Lin, Zhao Li, Sheng Zhou, Shichang Hu, Jialun Zhang, Linhao Luo, Jiarun Zhang, Longtao Huang, and Yuan He. 2022. Rmgn: A regional mask guided network for parser-free virtual try-on. *Proceedings of the Thirty-First International Conference on Artificial Intelligence* (2022).
- [32] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.
- [33] Matur Rahman Minar, Thai Thanh Tuan, Heejune Ahn, Paul Rosin, and Yu-Kun Lai. 2020. CP-VTON+: Clothing Shape and Texture Preserving Image-Based Virtual Try-On. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*.
- [34] Aymen Mir, Thiemo Alldieck, and Gerard Pons-Moll. 2020. Learning to transfer texture from clothing images to 3d humans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7023–7034.
- [35] Davide Morelli, Alberto Baldrati, Giuseppe Cartella, Marcella Cornia, Marco Bertini, and Rita Cucchiara. 2023. LaDI-VTON: latent diffusion textual-inversion enhanced virtual try-on. In *Proceedings of the 31st ACM International Conference on Multimedia*. 8580–8589.
- [36] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J Black. 2017. ClothCap: Seamless 4D Clothing Capture and Retargeting. *ACM Transactions on Graphics* 36, 4 (2017), 1–15.
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [38] Yurui Ren, Ge Li, Shan Liu, and Thomas H Li. 2020. Deep spatial transformation for pose-guided person image generation and animation. *IEEE Transactions on Image Processing* 29 (2020), 8622–8635.
- [39] Yurui Ren, Xiaoming Yu, Ruonan Zhang, Thomas H Li, Shan Liu, and Ge Li. 2019. StructureFlow: Image Inpainting via Structure-Aware Appearance Flow. In *Proceedings of the IEEE International Conference on Computer Vision*. 181–190.
- [40] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 10684–10695.
- [41] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. Springer, 234–241.
- [42] Igor Santesteban, Nils Thuerey, Miguel A Otaduy, and Dan Casas. 2021. Self-supervised collision handling via generative 3d garment models for virtual try-on. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 11763–11773.
- [43] Yoones A Sekhavat. 2016. Privacy Preserving Cloth Try-On Using Mobile Augmented Reality. *IEEE Transactions on Multimedia* 19, 5 (2016), 1041–1049.
- [44] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. *International Conference on Learning Representations* (2015), 1–14.
- [45] Bochao Wang, Huabin Zheng, Xiaodan Liang, Yimin Chen, Liang Lin, and Meng Yang. 2018. Toward Characteristic-Preserving Image-Based Virtual Try-On Network. In *Proceedings of the European Conference on Computer Vision*. 589–604.
- [46] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image Quality Assessment: from Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612.

- [47] Katja Wolff, Philipp Herholz, Verena Ziegler, Frauke Link, Nico Brügel, and Olga Sorkine-Hornung. 2022. Designing Personalized Garments with Body Movement. In *Computer Graphics Forum*. Wiley Online Library.
- [48] Zhonghua Wu, Guosheng Lin, Qingyi Tao, and Jianfei Cai. 2019. M2e-try on net: Fashion from model to everyone. In *Proceedings of the 27th ACM International Conference on Multimedia*. 293–301.
- [49] Zhenyu Xie, Zaiyu Huang, Xin Dong, Fuwei Zhao, Haoye Dong, Xijin Zhang, Feida Zhu, and Xiaodan Liang. 2023. Gp-vton: Towards general purpose virtual try-on via collaborative local-flow global-parsing learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 23550–23559.
- [50] Zhenyu Xie, Zaiyu Huang, Fuwei Zhao, Haoye Dong, Michael Kampffmeyer, and Xiaodan Liang. 2021. Towards Scalable Unpaired Virtual Try-On via Patch-Routed Spatially-Adaptive GAN. *Advances in Neural Information Processing Systems* 34 (2021), 2598–2610.
- [51] Zhenyu Xie, Xujie Zhang, Fuwei Zhao, Haoye Dong, Michael C Kampffmeyer, Haonan Yan, and Xiaodan Liang. 2021. Was-vton: Warping architecture search for virtual try-on network. In *Proceedings of the 29th ACM International Conference on Multimedia*. 3350–3359.
- [52] Rui Xu, Xiaoxiao Li, Bolei Zhou, and Chen Change Loy. 2019. Deep flow-guided video inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3723–3732.
- [53] Keyu Yan, Tingwei Gao, Hui Zhang, and Chengjun Xie. 2023. Linking Garment With Person via Semantically Associated Landmarks for Virtual Try-On. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 17194–17204.
- [54] Han Yang, Xinrui Yu, and Ziwei Liu. 2022. Full-range virtual try-on with recurrent tri-level transform. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3460–3469.
- [55] Han Yang, Ruimao Zhang, Xiaobao Guo, Wei Liu, Wangmeng Zuo, and Ping Luo. 2020. Towards Photo-Realistic Virtual Try-On by Adaptively Generating-Preserving Image Content. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7850–7859.
- [56] Zhijing Yang, Junyang Chen, Yukai Shi, Hao Li, Tianshui Chen, and Liang Lin. 2023. OccluMix: Towards De-Occlusion Virtual Try-on by Semantically-Guided Mixup. *IEEE Transactions on Multimedia* (2023).
- [57] Ruiyun Yu, Xiaoqi Wang, and Xiaohui Xie. 2019. Vtnfp: An image-based virtual try-on network with body and clothing feature preservation. In *Proceedings of the IEEE International Conference on Computer Vision*. 10511–10520.
- [58] Wei Yu, Yanping Li, Rui Wang, Wenming Cao, and Wei Xiang. 2022. PCFN: Progressive Cross-Modal Fusion Network for Human Pose Transfer. *IEEE Transactions on Circuits and Systems for Video Technology* (2022).
- [59] Shengping Zhang, Xiaoyu Han, Weigang Zhang, Xiangyuan Lan, Hongxun Yao, and Qingming Huang. 2023. Limb-Aware Virtual Try-On Network with Progressive Clothing Warping. *IEEE Transactions on Multimedia* (2023).
- [60] Xiaojing Zhong, Zhonghua Wu, Taizhe Tan, Guosheng Lin, and Qingyao Wu. 2021. Mv-ton: Memory-based video virtual try-on network. In *Proceedings of the 29th ACM International Conference on Multimedia*. 908–916.
- [61] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. 2016. View Synthesis by Appearance Flow. In *Proceedings of the European Conference on Computer Vision*. 286–301.
- [62] Luyang Zhu, Dawei Yang, Tyler Zhu, Fitsum Reda, William Chan, Chitwan Saharia, Mohammad Norouzi, and Ira Kemelmacher-Shlizerman. 2023. Tryondif-fusion: A tale of two unets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4606–4615.