

## GP-VTON: Towards General Purpose Virtual Try-on via Collaborative Local-Flow Global-Parsing Learning

Zhenyu Xie<sup>1</sup>, Zaiyu Huang<sup>1</sup>, Xin Dong<sup>2</sup>, Fuwei Zhao<sup>1</sup>  
 Haoye Dong<sup>3</sup>, Xijin Zhang<sup>2</sup>, Feida Zhu<sup>2</sup>, Xiaodan Liang<sup>\*1,4</sup>

<sup>1</sup>Shenzhen Campus of Sun Yat-Sen University, <sup>2</sup>ByteDance

<sup>3</sup>Carnegie Mellon University, <sup>4</sup>Peng Cheng Laboratory

{xiezy6, huangzy225, zhaofw}@mail2.sysu.edu.cn, donghaoye@cmu.edu

{dongxin.1016, zhangxijin}@bytedance.com, zhufeida@connect.hku.hk, xdliang328@gmail.com



Figure 1. Our method (A) outperforms existing SOTA methods (e.g., PF-AFN [12], FS-VTON [19]) on the challenging try-on scenario, and (B) can be easily extended to multi-category scenario to generate high-resolution photo-realistic try-on results.

### Abstract

Image-based Virtual Try-ON aims to transfer an in-shop garment onto a specific person. Existing methods employ a global warping module to model the anisotropic deformation for different garment parts, which fails to preserve the semantic information of different parts when receiving challenging inputs (e.g, intricate human poses, difficult garments). Moreover, most of them directly warp the input garment to align with the boundary of the preserved region, which usually requires texture squeezing to meet the boundary shape constraint and thus leads to texture distortion. The above inferior performance hinders existing methods from real-world applications. To address these problems and take a step towards real-world virtual try-on, we propose a General-Purpose Virtual Try-ON framework, named GP-VTON, by developing an innovative Local-Flow Global-Parsing (LFGP) warping module and a Dynamic

Gradient Truncation (DGT) training strategy. Specifically, compared with the previous global warping mechanism, LFGP employs local flows to warp garments parts individually, and assembles the local warped results via the global garment parsing, resulting in reasonable warped parts and a semantic-correct intact garment even with challenging inputs. On the other hand, our DGT training strategy dynamically truncates the gradient in the overlap area and the warped garment is no more required to meet the boundary constraint, which effectively avoids the texture squeezing problem. Furthermore, our GP-VTON can be easily extended to multi-category scenario and jointly trained by using data from different garment categories. Extensive experiments on two high-resolution benchmarks demonstrate our superiority over the existing state-of-the-art methods.<sup>1</sup>

<sup>1</sup>Corresponding author is Xiaodan Liang. Code is available at [gp-vton](https://github.com/xiaodanliang/gp-vton).

## 1. Introduction

The problem of Virtual Try-ON (VTON), aiming to transfer a garment onto a specific person, is of particular importance for the nowadays e-commerce and the future metaverse. Compared with the 3D-based solutions [2, 14, 16, 28, 34] which rely upon 3D scanning equipment or labor-intensive 3D annotations, the 2D image-based methods [1, 9, 12, 17, 19, 20, 29, 30, 32, 36, 38–40, 43], which directly manipulate on the images, are more practical for the real world scenarios and thus have been intensively explored in the past few years.

Although the pioneering 2D image-based VTON methods [12, 19, 29] can synthesize compelling results on the widely used benchmarks [6, 8, 18], there still exist some deficiencies preventing them from the real-world scenarios, which we argue mainly contain three-folds. First, existing methods have strict constraints on the input images, and are prone to generate artifacts when receiving challenging inputs. To be specific, as shown in the 1st row of Fig. 1(A), when the pose of the input person is intricate, existing methods [12, 19] fail to preserve the semantic information of different garment parts, resulting in the indistinguishable warped sleeves. Besides, as shown in the 2nd row of Fig. 1(A), if the input garment is a long sleeve without obvious seam between the sleeve and torso, existing methods will generate adhesive artifact between the sleeve and torso. Second, most of the existing methods directly squeeze the input garment to make it align with the preserved region, leading to the distorted texture around the preserved region (e.g., the 3rd row of Fig. 1(A)). Third, most of the existing works only focus on the upper-body try-on and neglect other garment categories (i.e., lower-body, dresses), which further limits their scalability for real-world scenarios.

To relieve the input constraint for VTON systems and fully exploit their application potential, in this paper, we take a step forwards and propose a unified framework, named GP-VTON, for the **General-Purposed Virtual Try-ON**, which can generate realistic try-on results even for the challenging scenario (Fig. 1(A)) (e.g., intricate human poses, difficult garment inputs, etc.), and can be easily extended to the multi-category scenario (Fig. 1(B)).

The innovations of our GP-VTON lie in a novel Local-Flow Global-Parsing (LFGP) warping module and a Dynamic Gradient Truncation (DGT) training strategy for the warping network, which enable the network to generate high fidelity deformed garments, and further facilitate our GP-VTON to generate photo-realistic try-on results.

Specifically, most of the existing methods employ neural network to model garment deformation by introducing the Thin Plate Splines (TPS) transformation [3] or the appearance flow [44] into the network, and training the network in a weakly supervised manner (i.e., without ground truth for the deformation function). However, both of the TPS-

based methods [6, 18, 30, 36, 40] and the flow-based methods [1, 12, 17, 19, 29] directly learn a global deformation field, therefore fail to represent complicated non-rigid garment deformation that requires diverse transformation for different garment parts. Taking the intricate pose case in Fig. 1(A) as an example, existing methods [12, 19] can not simultaneously guarantee accurate deformation for the torso region and sleeve region, and lead to exceeding distorted sleeves. In contrast, our LFGP warping module chooses to learn diverse local deformation fields for different garment parts, which is capable of individually warping each garment part, and generating semantic-correct warped garment even for intricate pose case. Besides, since each local deformation field merely affects one corresponding garment part, garment texture from other parts is agnostic to the current deformation field and will not appear in the current local warped result. Therefore, the garment adhesion problem in the complex garment scenario can be completely addressed (as demonstrated in the 2nd row of Fig. 1(A)). However, directly assembling the local warped parts together can not obtain realistic warped garments, because there would be overlap among different warped parts. To deal with this, our LFGP warping module collaboratively estimates a global garment parsing to fuse different local warped parts, resulting in a complete and unambiguous warped garment.

On the other hand, the warping network in existing methods [12, 19, 29] takes as inputs the flat garment and the mask of the preserved region (i.e., region to be preserved during the try-on procedure, such as the lower garment for upper-body VTON), and force the input garment to align with the boundary of the preserved region (e.g., the junction of the upper and lower garment), which usually require garment squeezing to meet the shape constraint and lead to texture distortion around the garment junction (please refer to the 3rd row of Fig. 1(A)). An effective solution to this problem is exploiting the gradient truncation strategy for the network training, in which the warped garment will be processed by the preserved mask before calculating the warping loss and the gradient in the preserved region will not be back-propagated. By using such a strategy, the warped garment is no longer required to strictly align with the preserved boundary, which largely avoids the garment squeezing and texture distortion. However, due to the poor supervision of warped garments in the preserved region, directly employing the gradient truncation for all training data will lead to excessive freedom for the deformation field, which usually results in texture stretching in the warped results. To tackle this problem, we proposed a Dynamic Gradient Truncation (DGT) training strategy which dynamically conducts gradient truncation for different training samples according to the disparity of height-width ratio between the flat garment and the warped garment. By introducing the dynamic mechanism, our LFGP warping module can alleviate

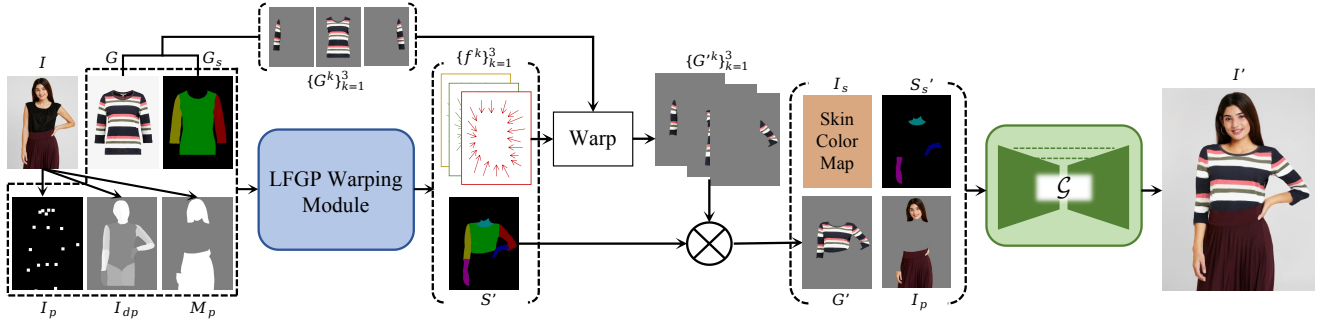


Figure 2. Overview of GP-VTON. The LFGP warping module aims to estimate the local flows  $\{f^k\}_{k=1}^3$  and global garment parsing  $S'$ , which is used to warp different garment parts  $\{G^k\}_{k=1}^3$  and assembles warped parts  $\{G'^k\}_{k=1}^3$  into the intact garment  $G'$ , respectively. The generator  $\mathcal{G}$  takes as inputs  $G'$  and the other person-related conditions to generate the final try-on result  $I'$ .

the texture stretching problem and obtain realistic warped garments with better texture preservation.

Overall, our contributions can be summarized as follows: (1) We propose a unified try-on framework, named GP-VTON, to generate photo-realistic results for diverse scenarios. (2) We propose a novel LFGP warping module to generate semantic-correct deformed garments even with challenging inputs. (3) We introduce a simple, yet effective DGT training strategy for the warping network to obtain distortion-free deformed garments. (4) Extensive experiments on two challenging high-resolution benchmarks show the superiority of GP-VTON over existing SOTAs.

## 2. Related Work

**Human-centric Image Synthesis.** Generative Adversarial Networks (GANs) [13], especially the StyleGAN-based models [24–27], have recently achieved significant success for photo-realistic image synthesis. Therefore, in the field of human synthesis, most of the existing methods [10, 11] inherit the StyleGAN-based architecture to obtain high-fidelity synthesized results. InsetGAN [10] combines the results from several pretrained GANs to a full-body human image, in which different pretrained GANs are in charge of the generation of different body parts (e.g., human body, face, hands, etc). StyleGAN-Human [11] explores three crucial factors for high-quality human synthesis, namely, dataset size, data distribution, and data alignment. In this paper, we focus on the image-based VTON, which aims to generate realistic human image via fitting an in-shop garment image onto a reference person.

**Image-based Virtual Try-on.** Most of the existing image-based VTON methods [1, 6–8, 12, 17, 19, 22, 29, 30, 32, 36, 40, 41] follow a two-stage generation framework that separately deforms the in-shop garment to target shape and synthesizes the try-on result via combining the deformed garment and the reference person. Since the quality of the garment deformation directly determines the realism of the generated results, it is crucial to design a powerful deformation module in this generation framework. Some previ-

ous methods [6, 8, 22, 36, 40, 41] leverage neural network to regress sparse garment control points in target image, which are then used to fit a TPS transformation [3] for garment deformation. Other methods [1, 7, 12, 17, 19] instead estimate an appearance flow map [44] to model non-rigid deformation, where the flow map depicts the corresponding location in the source image for each pixel in the target image. Compared with TPS-based methods which fit a transformation function via the sparse correspondence between control points, the flow-based methods directly predict the dense correspondence for each pixel, thus is more expressive for complex deformation. However, both of the existing TPS- and flow-based methods directly learn a global deformation field for various garment parts, which is unable to model diverse local transformation for different garment parts. Therefore, they fail to obtain realistic deformed results when receiving the intricate human pose. In this paper, we innovatively learn diverse local deformation fields for different garment parts, thus is capable of handling challenging inputs. Furthermore, existing methods usually neglect the texture distortion around the protected region. To address this problem, we propose a dynamic gradient truncation strategy for network training.

## 3. Methodology

Image-based virtual try-on algorithm aims to seamlessly transfer an in-shop garment  $G$  onto a specific person  $I$ . To achieve this, our GP-VTON proposes a Local-Flow Global-Parsing (LFGP) warping module (Sec. 3.1) to warp the garment to target shape, which first deforms local garment parts  $\{G^k\}_{k=1}^3$  individually and then assembles different warped parts  $\{G'^k\}_{k=1}^3$  together to obtain an intact warped garment  $G'$ . Besides, to address the texture distortion problem, GP-VTON introduces a Dynamic Gradient Truncation (DGT) training strategy (Sec. 3.2) for the warping network. Finally, GP-VTON employs a try-on generator (Sec. 3.3) to synthesize the try-on result  $I'$  according to  $G'$  and other person-related inputs. Furthermore, GP-

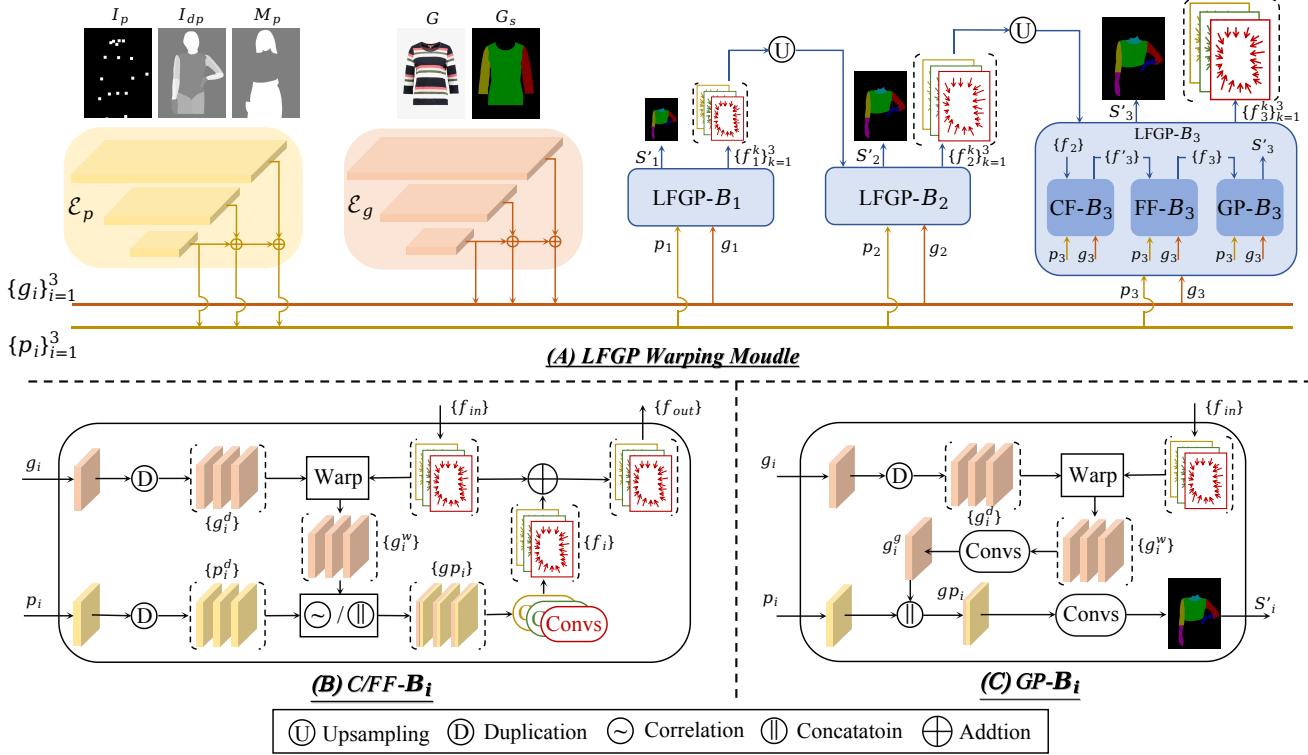


Figure 3. Overview of the LFGP warping module.

VTON can be easily extended for multi-category scenario and jointly trained by using data from various categories (Sec. 3.4). An overview of GP-VTON is displayed in Fig. 2.

### 3.1. Local-Flow Global-Parsing Warping Module

As shown in Fig. 3, our LFGP warping module follows the flow estimation pipeline in [12, 17, 19, 29] and is composed of pyramid feature extraction and cascade flow estimation. We will explain our dedicated improvement bellow.

**Pyramid Feature Extraction.** Our LFGP warping module employs two Feature Pyramid Network (FPN) [31] (i.e.,  $\mathcal{E}_p$  and  $\mathcal{E}_g$  in Fig. 3) to separately extract the multi-scale person feature  $\{p_i\}_{i=1}^N$  and garment feature  $\{g_i\}_{i=1}^N$ . Specifically,  $\mathcal{E}_p$  takes as inputs the human pose  $I_p$ , densepose pose  $I_{dp}$ , and the preserve region mask  $M_p$ , in which  $I_p$  and  $I_{dp}$  jointly provide the human pose information for flow estimation, and  $M_p$  is essential for the generation of the preserved-region-aware parsing.  $\mathcal{E}_g$  takes as inputs the intact in-shop garment  $G$  and its corresponding parsing map  $G_s$ , in which  $G_s$  can explicitly provide the semantic information of different garment parts for parsing generation. It is worth noting that, we extract five multi-scale features in our model (i.e.,  $N = 5$ ) but set  $N = 3$  in Fig. 3 for brevity.

#### Cascade Local-Flow Global-Parsing Estimation.

Most of the existing methods [12, 17, 19, 29] directly leverage a global flow to warp the intact garment, which tends to generate unrealistic warped result when different

garment parts require diverse deformation. To solve this problem, our LFGP module explicitly divides the intact garment into three local parts, (i.e., left/right sleeve, and torso region), and estimates three local flows to warp different parts individually. Since the deformation diversity within the same part is slight, the local flow can handle warping precisely and produce semantic-correct warped result. Furthermore, our LFGP estimates a global garment parsing to assemble local parts into an intact garment.

To be specific, as show in Fig. 3 (A), LFGP warping module exploits  $N$  LFGP blocks to cascadingly estimate  $N$  multi-scale local flows  $\{\{f_i^k\}_{k=1}^3\}_{i=1}^N$  and global garment parsing  $\{S'_i\}_{i=1}^N$ . Each LFGP block is composed of Coarse/Fine Flow Block (C/FF-B) and a Garment Parsing Block (GP-B), which estimate coarse/fine local flows and global garment parsing, respectively. As depicted in Fig. 3 (B), CF-B first duplicates the garment feature  $g_i$  and employs the incoming local flows  $\{f_{in}\}$ , which comes from previous LFGP block, to warp the duplicated garment feature  $\{g_i^d\}$  to three part-aware local warped features  $\{g_i^w\}$ . Then, the correlation operator from flownet2 [21] is employed to integrate the  $\{g_i^w\}$  and duplicated person feature  $\{p_i^d\}$  into three local fused features  $\{gp_i\}$ , which are separately sent into three convolution layers to estimate the corresponding local flows  $\{f'\}$ . At last,  $\{f'\}$  is added to  $\{f_{in}\}$  and produce the refined local flows  $\{f_{out}\}$ , which are the outputs of CF-B. FF-B has the same architecture as CF-F



Figure 4. Comparison of different training strategies (i.e., gradient truncation (GT) and dynamic gradient truncation (DGT)) for LFGP. In the red region, the warped result will be penalized and the gradient will be backpropagated. In the green region, the warped result will be neglected and the gradient will be truncated.

except it regards the output of CF-B as  $\{f_{in}\}$  and directly concatenate  $\{g_i^w\}$  and  $\{p_i^d\}$  to obtain  $\{gp_i\}$ . For GP-B, as shown in Fig. 3 (C), it employs the refined local flows  $\{f_{in}\}$  from FF-B to warp the duplicated feature  $\{g_i^d\}$  to the part-aware local features  $\{g_i^w\}$ , which are then fused by convolution layers and becomes a global warped feature  $g_i^g$ . Finally, the concatenation of  $g_i^g$  and the incoming  $p_i$  is passed to convolution layers to estimate the global garment parsing  $S'_i$ , whose labels consist of background, left/right sleeve, torso, left/right arms, and neck. With the strong guidance of the warped feature  $g_i^g$ , GP-B is prone to generate garment parsing that the garment shape in different local region is consistent with its corresponding local warped part.

After finishing the last estimation in LFGP module, as displayed in Fig. 2, GP-VTON deforms the local parts  $\{G^k\}_{k=1}^3$  individually via their corresponding local flows  $\{f^k\}_{k=1}^3$ , and assembles local warped parts to an intact warped garment  $G'$  by using the global garment parsing  $S'$ .

It is worth noting that, the global garment parsing is crucial for our local warping mechanism. Since there will be overlap among different warped parts, directly assembling the warped parts together will lead to distinct artifact in the overlap region. Instead, with the guidance of the global garment parsing, each pixel in the intact warped garment should be derived from a particular warped part, therefore the overlap artifact can be completely eliminated.

### 3.2. Dynamic Gradient Truncation

Existing methods [12, 19, 29, 36, 40] warp the in-shop garment according to the preserved region mask and force the warped garment to align with the boundary of the preserved region. However, directly warping the garment to meet the boundary constraint will lead to texture squeezing around the preserved region when the input person is in tucking-in style (as shown in the 1st case of Fig. 4.)

An intuitive solution for this issue is using the preserved mask to process the warped garment before calculating the training loss. In this way, the gradient in the preserve re-




Figure 5. Illustration of the garment partition and the ground truth of garment parsing for different categories.

gion will be truncated and the warped garment is no longer required to align with the boundary. However, when the training data is in tucking-out style (as shown in the 2nd case of Fig. 4.), gradient truncation is not appropriate since the inaccurate warped result in preserved region will not be penalized, leading to a stretched warped result.

To address above problems, our DGT training strategy dynamically conducts gradient truncation for different training samples according to their wearing style (i.e., tucking-in or tucking-out), in which the wearing style is determined by the disparity of the height-width ratio between the flat garment and the real warped garment (extracted from the person image). Fig. 4 provides an intuitive comparison among different training strategies. Specifically, we first extract the bounding boxes for the torso region of the flat garment and the warped garment by using their corresponding garment parsing. Then, we separately calculate the height-width ratio for each bounding box and use the ratio  $R_{style}$  between the warped garment item  $R^{warped}$  and flat garment item  $R^{flat}$  to reflect the wearing style of current training sample, which can be formulated as:

$$R_{style} = R^{warped} / R^{flat}, \quad (1)$$

where  $R^* = H^* / W^*$ , and  $H^*$  and  $W^*$  represent the height and width of the bounding box, respectively. We empirically find that if the person image is in tucking-in style,  $R_{style}$  is usually less than 0.9, and if it is in tucking-out style,  $R_{style}$  is usually more than 0.95. Therefore, we adopt gradient truncation when  $R_{style}$  less than 0.9 and abolish it when  $R_{style}$  more than 0.95. For training sample with  $R_{style}$  between 0.9 and 0.95, we randomly adopt gradient truncation with a probability of 0.5.

### 3.3. Try-on Generator

After the flow-parsing estimation stage, GP-VTON employs a Res-UNet-based [35] generator  $\mathcal{G}$  to synthesize the try-on result  $I'$ . As shown in Fig. 2,  $\mathcal{G}$  takes as inputs a skin color map  $I_s$ , a skin parsing map  $S'_s$ , a warped garment  $G'$ , and an image of preserved region  $I_p$ , in which  $I_s$  is a three-channel RGB image with the median value of the skin region (i.e., face, neck, arms.) while  $S'_s$  is a one-channel label map that contains the skin region (i.e., neck, arms) in  $S'$ .

### 3.4. Multi-category Virtual Try-on

Our GP-VTON can be easily extended to multi-category scenario through slight modifications. The core idea is using a unified partition mechanism for various garment categories. As mentioned in Dresscode [32], common garments can be classified into three macro-categories (i.e., upper, lower, and dresses). Since upper garment and dresses have the similar topology, we can apply the same partition mechanism for upper garment and dresses, namely, dividing the garment into the left/right part (i.e., left/right sleeve) and the middle part (i.e., torso region). To make lower garment consistent with the other categories, we regard the pants and skirt as a single type and divide it into three parts, which is also composed of left/right-part (i.e., left/right pant leg), and middle-part (i.e., skirt). By using this partition mechanism, garment from arbitrary categories can be divided into three local parts, which will be deformed individually and assembled to intact warped garment by our LFQP warping module. Besides, in the multi-category scenario, the estimated garment parsing is extended to include the labels of lower garment, which contain left/right pants, and the skirt. Fig. 5 displays the garment partition and the ground truth garment parsing for different garment categories.

### 3.5. Object Functions

During training, we train LFQP warping module and the generator separately. For LFQP, we utilize  $l_1$  loss  $\mathcal{L}_1$  and perceptual loss [23]  $\mathcal{L}_{per}$  for the warped results, and use  $l_1$  loss  $\mathcal{L}_m$  for the warped masks. We also use the pixel-wise cross-entropy loss  $\mathcal{L}_{ce}$  and the adversarial loss  $\mathcal{L}_{adv}$  for the estimated parsing. Besides, we follow PFAFN [12] and employ the second-order smooth loss  $\mathcal{L}_{sec}$  for the estimated flow. The total loss for LFQP module can be formulated as:

$$\mathcal{L}^w = \mathcal{L}_1^w + \lambda_{per}^w \mathcal{L}_{per}^w + \lambda_m^w \mathcal{L}_m^w + \lambda_{ce} \mathcal{L}_{ce} + \lambda_{adv}^w \mathcal{L}_{adv}^w + \lambda_{sec} \mathcal{L}_{sec}. \quad (2)$$

For the generator, we utilize  $l_1$  loss  $\mathcal{L}_1$ , the perceptual loss [23]  $\mathcal{L}_{per}$ , and the adversarial loss for the try-on result  $I'$ , and also utilize the  $l_1$  loss  $\mathcal{L}_m$  for the alpha mask  $M_c$ . The total loss is defined as follows:

$$\mathcal{L}^g = \mathcal{L}_1^g + \lambda_{per}^g \mathcal{L}_{per}^g + \lambda_{adv} \mathcal{L}_{adv} + \lambda_m^g \mathcal{L}_m^g. \quad (3)$$

More details are provided in the supplementary materials.

## 4. Experiments

**Datasets.** Our experiments are conducted under the resolution of  $512 \times 384$  by using two existing high-resolution virtual try-on benchmarks VITON-HD [6] and DressCode [32]. VITON-HD contains 13,679 image pairs of front-view upper-body woman and upper garment, which

Method	SSIM $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$	mIoU $\uparrow$	HE $\uparrow$
PF-AFN [12]	0.8858	9.475	0.0871	0.8412	14.9%
FS-VTON [19]	0.8829	9.552	0.0906	0.8357	8.80%
HR-VITON [29]	0.8623	16.21	0.1094	0.6949	9.10%
SDAFN [1]	0.8821	9.400	0.0922	0.5927	16.3%
<b>GP-VTON (Ours)</b>	<b>0.8939</b>	<b>9.197</b>	<b>0.0799</b>	<b>0.8764</b>	<b>50.9%</b>

Table 1. Quantitative comparisons on VITON-HD dataset [6]

are further split into 11,647/2,032 training/testing pairs. DressCode is composed of 48,392/5,400 training/testing pairs of front-view full-body person and garment from different categories (i.e., upper, lower, dresses). For each dataset, we employ [4] and [15] to extract the 2D pose and densepose, respectively. Besides, we apply a unified parsing estimator to predict the human/garment parsing for person/garment image, in which the estimator is based on [5] and trained by using 80k manual annotated fashion images.

**Baselines and Evaluation Metrics.** We compare GP-VTON with several state-of-the-art methods, including PF-AFN [12], FS-VTON [19], HR-VITON [29], and SDAFN [1], which are trained from scratch on VITON-HD [6] and DressCode [32] through using the official codes provided by the authors.

We employ three widely used metrics (i.e., Structural SIMilarity index (SSIM) [37], Perceptual distance (LPIPS) [42], and Fréchet Inception Distance (FID) [33]) to evaluate the similarity between synthesized and real images, in which SSIM and LPIPS are used for paired setting and FID are used for unpaired setting. We also utilize the mean Intersection over Union (mIoU) between the warped garment parsing and its corresponding ground truth (extracted from the human parsing) to evaluate the semantic-correctness of the warping module in different methods. Furthermore, we conduct Human Evaluation (HE) to evaluate different methods according to their synthesis quality.<sup>2</sup>

### 4.1. Qualitative Results

Fig. 6 and Fig. 7 display the qualitative comparison of GP-VTON with the state-of-the-art baselines on VITON-HD dataset [6] and DressCode dataset [32], respectively. Both figures demonstrate the superiority of GP-VTON over the baselines. First of all, the baselines fail to generate semantic-correct try-on results when encountering intricate poses, resulting in the damaged sleeves and arms (e.g., 1st row in Fig. 6), the blended pant legs and the indistinguishable sleeve (e.g., 1st case of 2nd row and 1st case of 3rd row in Fig. 7). Second, when receiving a complex garment (i.e., without obvious interval between adjacent parts), the baselines are prone to generate adhesive artifact (e.g., 2nd

<sup>2</sup>More details about architecture details, implementation details, and user study setting are provided in the supplementary materials.

Dataset	DressCode-Upper					DressCode-Lower					DressCode-Dresses				
	Method	SSIM $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$	mIoU $\uparrow$	HE $\uparrow$	SSIM $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$	mIoU $\uparrow$	HE $\uparrow$	SSIM $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$	mIoU $\uparrow$
PF-AFN [12]	0.9454	14.32	0.0380	0.8392	14.0%	0.9378	18.32	0.0445	0.9463	12.3%	0.8869	13.59	0.0758	0.8743	15.0%
FS-VTON [19]	0.9457	13.16	0.0376	0.8381	5.33%	0.9381	17.99	0.0438	0.9478	14.7%	<b>0.8876</b>	13.87	0.0745	0.8760	8.33%
HR-VITON [29]	0.9252	16.86	0.0635	0.6660	3.00%	0.9119	22.81	0.0811	0.8670	2.67%	0.8642	16.12	0.1132	0.7209	2.33%
SDAFN [1]	0.9379	12.61	0.0484	0.5046	11.3%	0.9317	<b>16.05</b>	0.0549	0.4543	13.3%	0.8776	<b>11.80</b>	0.0852	0.5945	19.3%
<b>GP-VTON (Ours)</b>	<b>0.9479</b>	<b>11.98</b>	<b>0.0359</b>	<b>0.8766</b>	<b>66.3%</b>	<b>0.9405</b>	16.07	<b>0.0420</b>	<b>0.9601</b>	<b>57.0%</b>	0.8866	12.26	<b>0.0729</b>	<b>0.8951</b>	<b>55.0%</b>

Table 2. Quantitative comparisons on DressCode dataset [32]



Figure 6. Qualitative comparison on VITON-HD dataset [6]. Please zoom in for more details.

Method	LF	GT	DGT	SSIM $\uparrow$	LPIPS $\downarrow$	mIoU $\uparrow$	$R_{diff} \downarrow$
LFGP $\dagger$	$\times$	$\times$	$\times$	0.9016	0.0950	0.8412	0.3058
LFGP $\star$	$\checkmark$	$\times$	$\times$	0.9039	0.0911	<b>0.8784</b>	0.3003
LFGP $\ast$	$\checkmark$	$\checkmark$	$\times$	<b>0.9053</b>	0.0900	0.8774	0.2409
LFGP	$\checkmark$	$\times$	$\checkmark$	0.9050	<b>0.0884</b>	0.8764	<b>0.1655</b>

Table 3. Ablation study of the Local FFlow (LF), Gradient Truncation (GT), and Dynamic Gradient Truncation (DGT) on the VITON-HD dataset [6].

row in Fig. 6 and 2nd case of 2nd row in Fig. 7). Third, existing methods [12, 19, 29] tend to generate distorted texture around the preserved region (e.g., the third row in Fig. 6). In comparison, GP-VTON first employs local flows to warp different garment parts individually, leading to the precise local warped parts, and then uses the global garment parsing to assemble local parts into an semantic-correct warped garment. Therefore, GP-VTON is more robust to intricate pose or complex input garment. Besides, by using the dynamic gradient truncation training strategy, GP-VTON can avoid generating distorted texture around preserved region.

## 4.2. Quantitative Results

As reported in Tab. 1, our GP-VTON consistently surpasses the baselines on all metrics for the VITON-HD dataset [6], demonstrating that GP-VTON can obtain more precise warped garments and generate try-on results with better visual quality. Particularly, on the mIoU metric, GP-VTON outperforms other methods by a large margin, which

further illustrates that our LFGP warping module is capable of obtaining the semantic-correct warped results. Tab. 2 shows the quantitative comparisons of GP-VTON with other methods on the DressCode dataset [32]. As shown in the table, for DressCode-Upper, GP-VTON achieves the finest score on all metrics. For DressCode-Lower and DressCode-Dresses, GP-VTON outperforms other methods on most of the metrics and obtains comparable low FID score with SDAFN [1]. This is mainly because the human pose in the DressCode-Lower and DressCode-Dresses are generally simple, which do not require complex warping during try-on process, thus the advanced SDAFN [1] can also obtain compelling FID score. However, the superiority of GP-VTON on the mIoU and HE metrics can still indicate its warped results are more semantically correct and its synthesized results are more photo-realistic.

## 4.3. Ablation Study

To validate the effectiveness of LFGP warping module and DGT training strategy, we design three variants of our proposed method and evaluate the performance of different variants according to their metric scores for the warped results. Besides, we define another metric  $R_{diff}$  to measure the difference of the height-width ratio between the warped garment and the in-shop garment, where the lower value indicates the better preservation of the original height-width ratio, thus implying the better warping.

We regard PF-AFN [12] as our first variant (denoted as LFGP $\dagger$ ), since it utilizes a global flow for warping and is



Figure 7. Qualitative comparison on Dresscode dataset [32]. Please zoom in for more details.

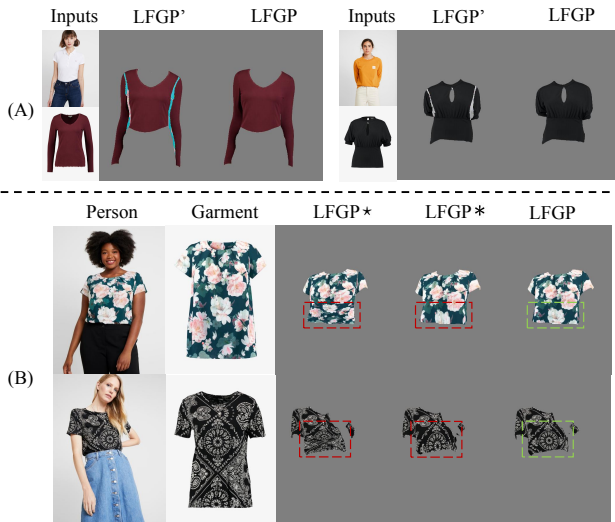


Figure 8. Ablation studies on the effectiveness of (A) the global parsing during the parts assembling process and (B) the dynamic gradient truncation training strategy.

trained without Gradient Truncation. We further implement the other two variants (i.e., LFGP $\star$  and LFGP $\star$ ) by training LFGP module without Gradient Truncation and without Dynamic Gradient Truncation, respectively.

**LFGP Module.** As reported in Tab. 3, compared with LFGP $\dagger$ , other methods with local flows gain increase on SSIM and LPIPS, and achieve obvious improvement on the mIoU metric, demonstrating that our local flows warping mechanism can obtain more realistic and semantic-correct warped results. Besides, we further conduct another experiment on the full LFGP model (denoted as LFGP $\dagger$ ) to demonstrate the effectiveness of the global parsing. As shown in Fig. 8 (A), by using the global parsing to assemble different warped parts, the overlap artifact between different parts can be completely eliminated.

**DGT Training Strategy.** As reported in Tab. 3, compared with LFGP $\star$ , LFGP $\star$  with the normal GT strategy obtains lower  $R_{diff}$  score while the full LFGP model with DGT

strategy achieves the lowest  $R_{diff}$  score. Fig. 8 (B) further provide the visual comparisons among different methods, in which LFGP $\star$  tends to squeeze the texture while LFGP $\dagger$  tends to stretch the texture. In contrast, the full LFGP module can preserve the texture details well. Both of the quantitative and qualitative comparisons validate that training with DGT facilitates the warping model to preserve the original height-width ratio of the garment, thus avoid texture squeezing or stretching.

## 5. Conclusion

In this work, we propose GP-VTON towards the general-purpose virtual try-on, which is capable of generating semantic-correct and photo-realistic try-on results even in the challenging self-occlusion scenarios and can be easily extended to multi-category scenarios. Specifically, to make garment warping robust to intricate inputs, GP-VTON introduces the Local-Flow Global-Parsing (LFGP) warping module to warp local parts individually and assembles local warped parts via the estimated global garment parsing. Besides, to alleviate the texture distortion problem in existing methods, GP-VTON employs a Dynamic Gradient Truncation (DGT) training strategy for the warping network. Experiments on two high resolution virtual try-on benchmark illustrate GP-VTON’s superiority over existing methods. The limitation and social impact of our GP-VTON will be discussed in the supplementary materials.

## 6. Acknowledgement

This work was supported in part by National Key R&D Program of China under Grant No.2020AAA0109700, National Natural Science Foundation of China (NSFC) under Grant No.61976233, Guangdong Outstanding Youth Fund (Grant No.2021B1515020061), Shenzhen Fundamental Research Program (Project No.RCYX20200714114642083, No.JCYJ20190807154211365), the Fundamental Research Funds for the Central Universities, Sun Yat-sen University under Grant No.221gqb38.



## References

- [1] Shuai Bai, Huiling Zhou, Zhikang Li, Chang Zhou, and Hongxia Yang. Single stage virtual try-on via deformable attention flows. In *ECCV*, 2022. 2, 3, 6, 7
- [2] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *ICCV*, pages 5420–5430, 2019. 2
- [3] F.L. Bookstein. Principal warps: thin-plate splines and the decomposition of deformations. *TPAMI*, 11(6):567–585, 1989. 2, 3
- [4] Z. Cao, G. Hidalgo, T. Simon, S. Wei, and Y. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *TPAMI*, 43(01):172–186, 2021. 6
- [5] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, pages 1290–1299, 2022. 6
- [6] Seunghwan Choi, Sunghyun Park, Minsoo Lee, and Jaegul Choo. Viton-hd: High-resolution virtual try-on via misalignment-aware normalization. In *CVPR*, pages 14131–14140, 2021. 2, 3, 6, 7
- [7] Ayush Chopra, Rishabh Jain, Mayur Hemani, and Balaji Krishnamurthy. Zflow: Gated appearance flow-based virtual try-on with 3d priors. In *ICCV*, pages 5433–5442, 2021. 3
- [8] Haoye Dong, Xiaodan Liang, Xiaohui Shen, Bochao Wang, Hanjiang Lai, Jia Zhu, Zhiting Hu, and Jian Yin. Towards multi-pose guided virtual try-on network. In *ICCV*, pages 9026–9035, 2019. 2, 3
- [9] Xin Dong, Fuwei Zhao, Zhenyu Xie, Xijin Zhang, Kang Du, Min Zheng, Xiang Long, Xiaodan Liang, and Jianchao Yang. Dressing in the wild by watching dance videos. In *CVPR*, pages 3480–3489, 2022. 2
- [10] Anna Frühstück, Krishna Kumar Singh, Eli Shechtman, Niloy J. Mitra, Peter Wonka, Jingwan Lu, Anna Frühstück, Krishna Kumar Singh, Eli Shechtman, Niloy J. Mitra, Peter Wonka, and Jingwan Lu. Insetgan for full-body image generation. In *CVPR*, pages 7723–7732, 2022. 3
- [11] Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen Change Loy, Wayne Wu, and Ziwei Liu. Stylegan-human: A data-centric odyssey of human generation. In *ECCV*, 2022. 3
- [12] Yuying Ge, Yibing Song, Ruimao Zhang, Chongjian Ge, Wei Liu, and Ping Luo. Parser-free virtual try-on via distilling appearance flows. In *CVPR*, pages 8485–8493, 2021. 1, 2, 3, 4, 5, 6, 7
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 3
- [14] Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. Drape: Dressing any person. *ACM Transactions on Graphics*, 31(4), 2012. 2
- [15] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *CVPR*, pages 7297–7306, 2018. 6
- [16] Fabian Hahn, Bernhard Thomaszewski, Stelian Coros, Robert W Sumner, Forrester Cole, Mark Meyer, Tony DeRose, and Markus Gross. Subspace clothing simulation using adaptive bases. *ACM Transactions on Graphics*, 33(4), 2014. 2
- [17] Xintong Han, Xiaojun Hu, Weilin Huang, and Matthew R. Scott. Clothflow: A flow-based model for clothed person generation. In *ICCV*, pages 10471–10480, 2019. 2, 3, 4
- [18] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S Davis. Viton: An image-based virtual try-on network. In *CVPR*, pages 7543–7552, 2018. 2
- [19] Sen He, Yi-Zhe Song, and Tao Xiang. Style-based global appearance flow for virtual try-on. In *CVPR*, pages 3470–3479, 2022. 1, 2, 3, 4, 5, 6, 7
- [20] Zaiyu Huang, Hanhui Li, Zhenyu Xie, Michael Kampffmeyer, Qingling Cai, and Xiaodan Liang. Towards hard-pose virtual try-on via 3d-aware global correspondence learning. In *NeurIPS*, 2022. 2
- [21] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017. 4
- [22] Thibaut Issenhuth, Jérémie Mary, and Clément Calauzènes. Do not mask what you do not need to mask: A parser-free virtual try-on. In *ECCV*, pages 619–635, 2020. 3
- [23] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711, 2016. 6
- [24] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *NeurIPS*, pages 12104–12114, 2020. 3
- [25] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *NeurIPS*, pages 852–863, 2021. 3
- [26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial network. In *ICCV*, pages 4401–4410, 2019. 3
- [27] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, pages 8110–8119, 2020. 3
- [28] Zorah Lahner, Daniel Cremers, and Tony Tung. Deepwrinkles: Accurate and realistic clothing modeling. In *ECCV*, pages 667–684, 2018. 2
- [29] Sangyun Lee, Gyojung Gu, Sunghyun Park, Seunghwan Choi, and Jaegul Choo. High-resolution virtual try-on with misalignment and occlusion-handled conditions. In *ECCV*, 2022. 2, 3, 4, 5, 6, 7
- [30] Kedan Li, Min Jin Chong, Jeffrey Zhang, and Jingen Liu. Toward accurate and realistic outfits visualization with attention to details. In *CVPR*, pages 15546–15555, 2021. 2, 3
- [31] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. 4
- [32] Davide Morelli, Matteo Fincato, Marcella Cornia, Federico Landi, Fabio Cesari, and Rita Cucchiara. Dress Code: High-

- Resolution Multi-Category Virtual Try-On. In *ECCV*, 2022. [2](#), [3](#), [6](#), [7](#), [8](#)
- [33] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *CVPR*, pages 11410–11417, 2022. [6](#)
- [34] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J Black. Clothcap: seamless 4d clothing capture and retargeting. *ACM Transactions on Graphics*, 36(4), 2017. [2](#)
- [35] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. [5](#)
- [36] Bochao Wang, Huabin Zheng, Xiaodan Liang, Yimin Chen, Liang Lin, and Meng Yang. Toward characteristic-preserving image-based virtual try-on network. In *ECCV*, pages 589–604, 2018. [2](#), [3](#), [5](#)
- [37] Zhou Wang, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. [6](#)
- [38] Zhenyu Xie, Zaiyu Huang, Fuwei Zhao, Haoye Dong, Michael Kampffmeyer, and Xiaodan Liang. Towards scalable unpaired virtual try-on via patch-routed spatially-adaptive gan. In *NeurIPS*, 2021. [2](#)
- [39] Zhenyu Xie, Xujie Zhang, Fuwei Zhao, Haoye Dong, Michael C Kampffmeyer, Haonan Yan, and Xiaodan Liang. Was-vton: Warping architecture search for virtual try-on network. In *ACMMM*, pages 3350–3359, 2021. [2](#)
- [40] Han Yang, Ruimao Zhang, Xiaobao Guo, Wei Liu, Wangmeng Zuo, and Ping Luo. Towards photo-realistic virtual try-on by adaptively generating-preserving image content. In *CVPR*, pages 7850–7859, 2020. [2](#), [3](#), [5](#)
- [41] Ruiyun Yu, Xiaoqi Wang, and Xiaohui Xie. Vtnfp: An image-based virtual try-on network with body and clothing feature preservation. In *ICCV*, pages 10511–10520, 2019. [3](#)
- [42] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [6](#)
- [43] Fuwei Zhao, Zhenyu Xie, Michael Kampffmeyer, Haoye Dong, Songfang Han, Tianxiang Zheng, Tao Zhang, and Xiaodan Liang. M3d-vton: A monocular-to-3d virtual try-on network. In *ICCV*, pages 13239–13249, 2021. [2](#)
- [44] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A. Efros. View synthesis by appearance flow. In *ECCV*, 2016. [2](#), [3](#)