

Supplementary Material for USC-PFN

Chenghu Du¹ Junyin Wang¹ Shuqin Liu⁴ Shengwu Xiong^{1,2,3,5,*}
¹Wuhan University of Technology, ²Shanghai AI Laboratory
³Sanya Science and Education Innovation Park, Wuhan University of Technology
⁴Wuhan Textile University, ⁵Qiongtai Normal University
 {duch, wjy199708, xiongsw}@whut.edu.cn, liushuqingwtu@foxmail.com
<https://du-chenghu.github.io/USC-PFN/>

1 Introduction

This supplementary material includes: 1) architecture details of USC-PFN, providing a comprehensive understanding of its methodology and development processes; 2) training and inference details of USC-PFN, providing a comprehensive understanding of its training and inference processes; and 3) additional qualitative results from different models, offering further insights and examples related to the research topic.

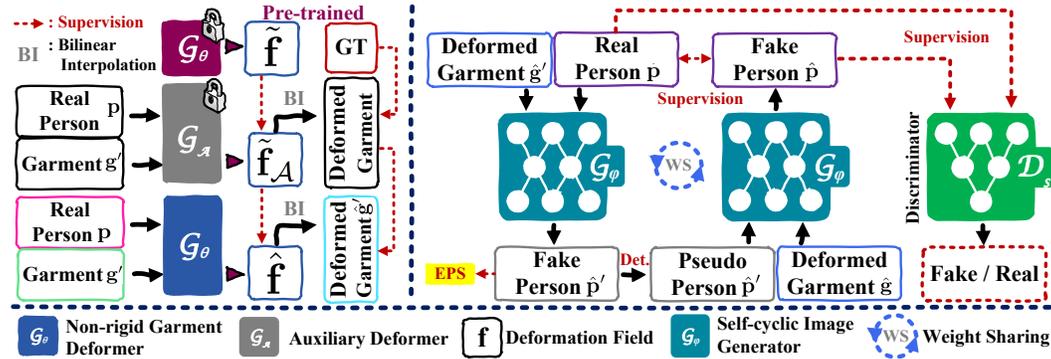


Figure 1: Overview of our self-cycle consistency network: paired garment-person images (g, p) and an arbitrary garment g' serve as the training data. **Left:** G_θ is pre-trained using paired images to obtain the ground-truth deformation field \tilde{f} for the auxiliary deformer G_A to optimize \tilde{f}_A . Then, G_θ is further trained by using \tilde{f}_A of G_A^* with unpaired images. **Right:** G_ϕ takes unpaired (g', p) to generate a fake person \hat{p}' . Then, the real person p is reconstructed by re-feeding \hat{p}' and \hat{g} into G_ϕ , ensuring self-cycle consistency $p \approx \hat{p}$. D_s is the discriminator that learns to distinguish real and fake images. **EPS** represents enhanced pixel-level supervision and **Det.** represents detaching tensor \hat{p}' from the computation graph.

2 Architecture Details

In our architecture, any encoder-decoder network can be used as our generator. In our experiments, to achieve fast convergence of the network during the theoretical validation phase, we adopt Res-UNet [1] as Non-rigid Garment Deformer G_θ (NGD), auxiliary deformer G_A , and Self-cyclic Image Generator G_ϕ (SIG). The discriminator is from pix2pixHD [2]. The complete code for Res-UNet can be found in our open-source repository <https://github.com/du-chenghu/USC-PFN/>.

*Shengwu Xiong is corresponding author.

3 Training and Inference Details

As shown in Figure 1, our framework consists of three generators (\mathcal{G}_θ , \mathcal{G}_A , and \mathcal{G}_φ) and one discriminator (\mathcal{D}_s).

3.1 Pre-trained Parser for Supervision

The first thing that needs to be specifically pointed out is that "**Parser-Free Virtual Try-On**" refers to the training and inference stages, where in the input of the main generator, there is no need for human parsing (i.e., using a parser-generated human semantic segmentation map). Because [3] has confirmed that including human parsing in the generator’s input, and if there are errors in this parsing, can negatively impact the results, eliminating human parsing (i.e. the parser) at the input end of the generator, is highly beneficial.

In order to facilitate the accurate localization of different semantic parts of the human body during the calculation of loss in the training stage, we utilized a pre-trained parser (denotes as τ). This parser, similar to the one used in [4, 5, 3, 6–8], is designed to predict a posterior "after-try-on" semantics, which corresponds to the semantic map of the desired try-on result. Only by inputting the person representation \mathcal{R} and the target garment \mathbf{g}' into τ , the target semantic map $\delta' \in \{0, 1\}^{21 \times H \times W}$ can be generated, which can be represented as:

$$\delta' = \tau(\mathcal{R}, \mathbf{g}'). \quad (1)$$

In the main text, we generate the result $\hat{\mathbf{p}}'$ using unpaired $(\mathbf{p}, \hat{\mathbf{g}}')$. With the help of δ' , $\hat{\mathbf{p}}'_s$ can be obtained through:

$$\hat{\mathbf{p}}'_s = (\delta'_{arms} \cup \delta'_{neck}) \odot \hat{\mathbf{p}}', \quad (2)$$

where \odot denotes entry-wise multiplication. $\hat{\mathbf{p}}'_g$ can be obtained through:

$$\hat{\mathbf{p}}'_g = \delta'_{garment} \odot \hat{\mathbf{p}}'. \quad (3)$$

$\hat{\mathbf{p}}'_c$ can be obtained through:

$$\delta'_{identity} = 1 - (\delta'_{arms} \cup \delta'_{neck} \cup \delta'_{garment}), \quad (4)$$

$$\hat{\mathbf{p}}'_c = \delta'_{identity} \odot \hat{\mathbf{p}}', \quad (5)$$

where $-$ denotes entry-wise subtraction and 1 is all-ones tensor.

3.2 Non-rigid Garment Deformer \mathcal{G}_θ (NGD)

Training. During the training phase, we first train \mathcal{G}_θ based on Markov Random Field theory using paired garment-person images:

$$\tilde{\mathbf{f}} = \mathcal{G}_\theta(\mathbf{p}, \mathbf{g}). \quad (6)$$

The reason for using paired (\mathbf{p}, \mathbf{g}) for training is that the clothes on \mathbf{p} and \mathbf{g} look the same in appearance. Therefore, we can learn implicit correlations between the two during the deformation process. These correlations could be related to color, shape, perspective changes, texture variations, shadows and lighting, depth of field, etc. It is conceivable that if \mathbf{p} is replaced with the person representation \mathcal{R} , no relevant information other than the shape can be extracted at all.

Unfortunately, \mathcal{G}_θ does not need correlations in terms of color, shape, and texture. This results in a significant entanglement of features between target clothing and clothing worn on the person in \mathcal{G}_θ ’s feature space. Consequently, \mathcal{G}_θ can only be effective when the input data is paired, and it fails to deform properly when encountering unpaired samples.

We employ an auxiliary deformer, \mathcal{G}_A , to address this issue. Since \mathcal{R} does not contain the aforementioned correlations except for the shape, we modify \mathcal{G}_A ’s input from (\mathbf{p}, \mathbf{g}) to $(\mathcal{R}, \mathbf{g})$ to estimate the target deformation field $\tilde{\mathbf{f}}_A$:

$$\tilde{\mathbf{f}}_A = \mathcal{G}_A(\mathcal{R}, \mathbf{g}). \quad (7)$$

However, the limited amount of information in \mathcal{R} forces us to utilize previously pre-trained \mathcal{G}_θ to supplement the aforementioned correlations for \mathcal{G}_A . In other words, we optimize \mathcal{G}_A ’s result $\tilde{\mathbf{f}}_A$ by leveraging the result $\tilde{\mathbf{f}}$ from \mathcal{G}_θ , minimizing the differences between $\tilde{\mathbf{f}}$ and $\tilde{\mathbf{f}}_A$:

$$\tilde{\mathbf{f}}_A^* = \operatorname{argmin}_{\tilde{\mathbf{f}}_A} \left\| \tilde{\mathbf{f}}_A - \tilde{\mathbf{f}} \right\|. \quad (8)$$

After obtaining the optimal $\mathcal{G}_{\mathcal{A}}^*$, we continue to train the deformer \mathcal{G}_{θ} by taking unpaired $(\mathbf{p}, \mathbf{g}')$ as input and $\mathbf{f}_{\mathcal{A}}$ as supervision, to disentangle the feature entanglement of \mathcal{G}_{θ} :

$$\tilde{\mathbf{f}}_{\mathcal{A}} \approx \hat{\mathbf{f}}, \text{ with } \hat{\mathbf{f}} = \mathcal{G}_{\theta}(\mathbf{p}, \mathbf{g}'), \tilde{\mathbf{f}}_{\mathcal{A}} = \mathcal{G}_{\mathcal{A}}^*(\mathcal{R}, \mathbf{g}'). \quad (9)$$

Inference. During the inference phase, we exclusively utilize \mathcal{G}_{θ} to warp arbitrary garment.

3.3 Self-cyclic Image Generator \mathcal{G}_{φ} (SIG)

Training. During the training phase, we divided the training of \mathcal{G}_{φ} into two steps as shown in Figure 1. In the first step, we train \mathcal{G}_{φ} taking unpaired $(\mathbf{p}, \hat{\mathbf{g}}')$ as input to obtain the fake result $\hat{\mathbf{p}}'$:

$$\hat{\mathbf{p}}' = \mathcal{G}_{\varphi}(\mathbf{p}, \hat{\mathbf{g}}'). \quad (10)$$

Then, in the second step, $\hat{\mathbf{p}}'$ detached from the computation graph and $\hat{\mathbf{g}}$ are used as input of \mathcal{G}_{φ} to reconstruct the real person:

$$\mathbf{p} \approx \hat{\mathbf{p}}, \text{ with } \hat{\mathbf{p}} = \mathcal{G}_{\varphi}(\hat{\mathbf{p}}', \hat{\mathbf{g}}). \quad (11)$$

In order to accelerate the convergence of \mathcal{G}_{φ} , \mathbf{p} and $\hat{\mathbf{p}}'$ can be combined using τ and used as input in the second step:

$$\hat{\mathbf{p}}' = \delta_{save} \odot \mathbf{p} + (1 - \delta_{save}) \odot \hat{\mathbf{p}}', \quad (12)$$

where

$$\delta = \tau(\mathcal{R}, \mathbf{g}), \quad (13)$$

$$\delta_{identity} = 1 - (\delta_{arms} \cup \delta_{neck} \cup \delta_{garment}), \quad (14)$$

$$\delta_{save} = (\delta_{identity} \cap \delta'_{identity}) \cup (\delta_{arms} \cap \delta'_{arms}) \cup (\delta_{neck} \cap \delta'_{neck}). \quad (15)$$

Additionally, parts of the skin $\mathbf{s}_{\mathbf{p}}$ can be obtained through:

$$\mathbf{s}_{\mathbf{p}} = ((\delta_{arms} \cap \delta'_{arms}) \cup (\delta_{neck} \cap \delta'_{neck})) \odot \mathbf{p}. \quad (16)$$

Inference. During the inference phase, the warped garment image and the person image are fed to \mathcal{G}_{φ} to synthesize the try-on result.

4 More Qualitative Results

We show more qualitative comparisons of the try-on results among our proposed USC-PFN and other cutting-edge methods (CP-VTON+[9], ACGPN[4], DCTON[5], PF-AFN[3], and FS-VTON[7]) on VITON Zalando dataset [10]. As shown in Figure 2. overall, our model generates better try-on images.

References

- [1] Foivos I Diakogiannis, François Waldner, Peter Caccetta, and Chen Wu. Resunet-a: A deep learning framework for semantic segmentation of remotely sensed data. *ISPRS Journal of Photogrammetry and Remote Sensing*, 162:94–114, 2020.
- [2] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.
- [3] Yuying Ge, Yibing Song, Ruimao Zhang, Chongjian Ge, Wei Liu, and Ping Luo. Parser-free virtual try-on via distilling appearance flows. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8485–8493, 2021.
- [4] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7850–7859, 2020.
- [5] Chongjian Ge, Yibing Song, Yuying Ge, Han Yang, Wei Liu, and Ping Luo. Disentangled cycle consistency for highly-realistic virtual try-on. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16928–16937, 2021.

- [6] Ayush Chopra, Rishabh Jain, Mayur Hemani, and Balaji Krishnamurthy. Zflow: Gated appearance flow-based virtual try-on with 3d priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5433–5442, 2021.
- [7] Sen He, Yi-Zhe Song, and Tao Xiang. Style-based global appearance flow for virtual try-on. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3470–3479, 2022.
- [8] Han Yang, Xinrui Yu, and Ziwei Liu. Full-range virtual try-on with recurrent tri-level transform. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3460–3469, 2022.
- [9] Matiur Rahman Minar, Thai Thanh Tuan, Heejune Ahn, Paul Rosin, and Yu-Kun Lai. Cp-vton+: Clothing shape and texture preserving image-based virtual try-on. In *CVPR Workshops*, volume 3, pages 10–14, 2020.
- [10] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S Davis. Viton: An image-based virtual try-on network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7543–7552, 2018.

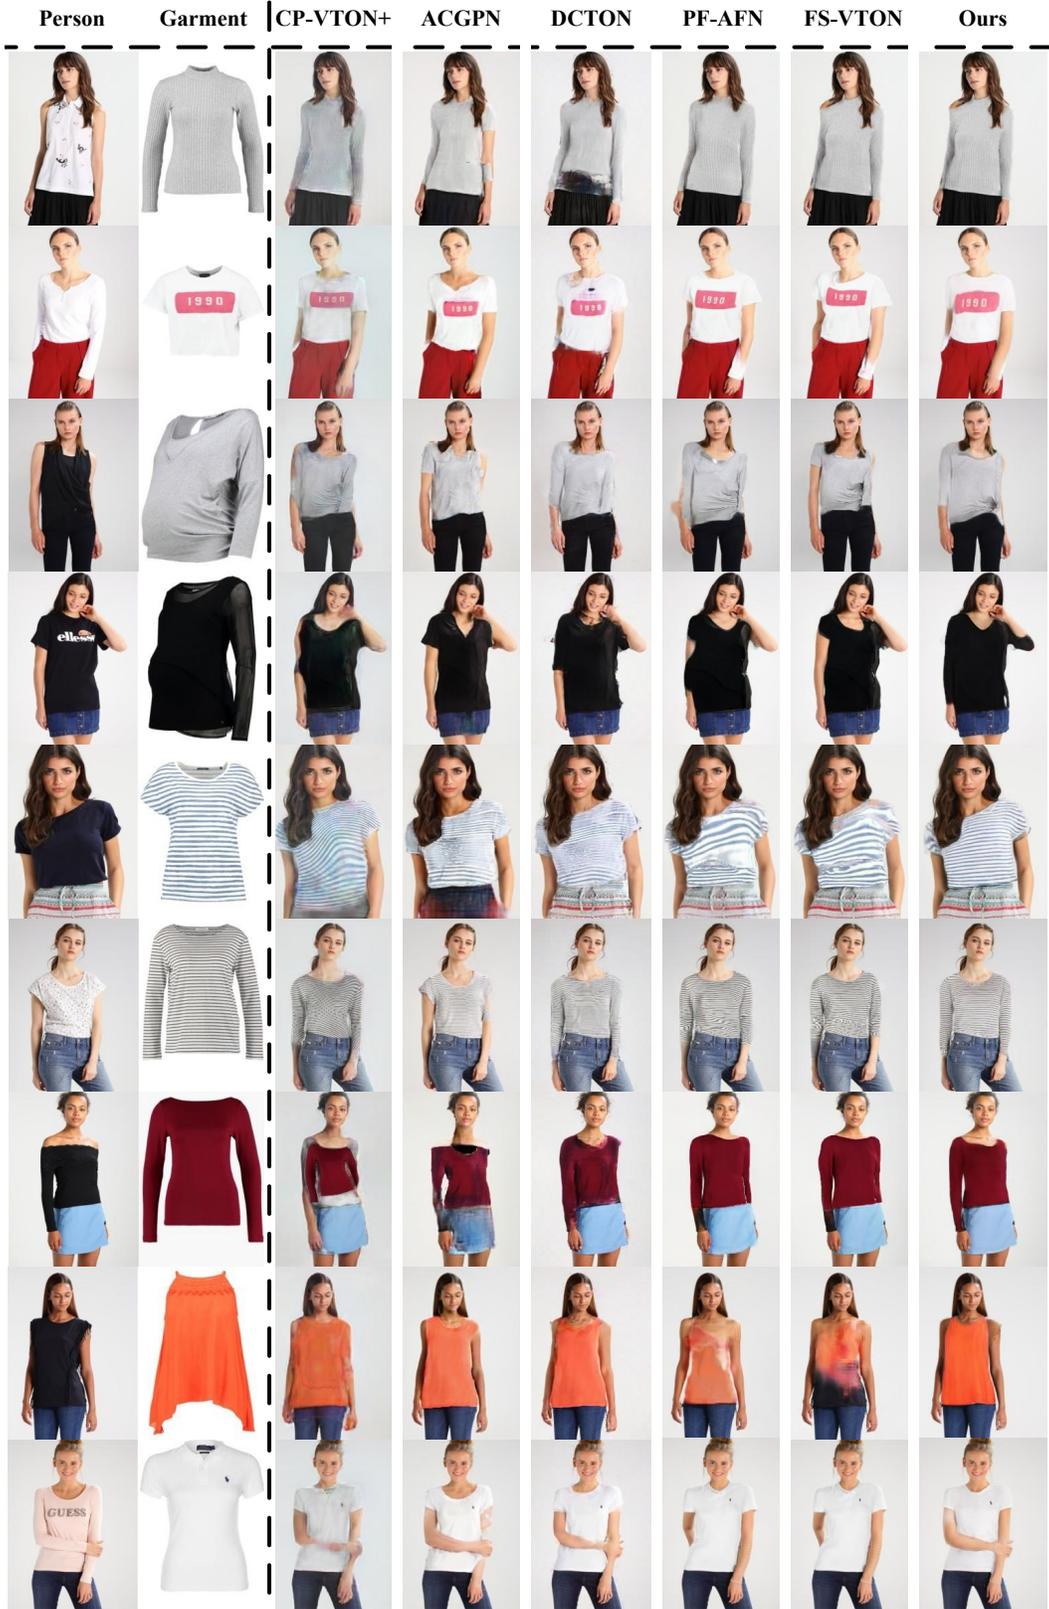


Figure 2: Qualitative results of different methods (CP-VTON+[9], ACGPN[4], DCTON[5], PF-AFN[3], FS-VTON[7], and ours) in the unpaired setting.