Fill in Fabrics: Body-Aware Self-Supervised Inpainting for Image-Based Virtual Try-On

Hasib Zunair¹ hasibzunair@gmail.com Yan Gobeil² yan.gobeil@decathlon.com Samuel Mercier² samuel.mercier@decathlon.com A. Ben Hamza¹

Montreal, QC, Canada ² Decathlon Canada Montreal, QC, Canada

¹ Concordia University

A. Ben Hamza¹ hamza@ciise.concordia.ca

Abstract

Previous virtual try-on methods usually focus on aligning a clothing item with a person, limiting their ability to exploit the complex pose, shape and skin color of the person, as well as the overall structure of the clothing, which is vital to photo-realistic virtual try-on. To address this potential weakness, we propose a fill in fabrics (FIFA) model, a self-supervised conditional generative adversarial network based framework comprised of a Fabricator and a unified virtual try-on pipeline with a Segmenter, Warper and Fuser. The Fabricator aims to reconstruct the clothing image when provided with a masked clothing as input, and learns the overall structure of the clothing by filling in fabrics. A virtual try-on pipeline is then trained by transferring the learned representations from the Fabricator to Warper in an effort to warp and refine the target clothing. We also propose to use a multi-scale structural constraint to enforce global context at multiple scales while warping the target clothing to better fit the pose and shape of the person. Extensive experiments demonstrate that our FIFA model achieves state-of-the-art results on the standard VITON dataset for virtual try-on of clothing items, and is shown to be effective at handling complex poses and retaining the texture and embroidery of the clothing.

1 Introduction

The core objective of image-based virtual try-on is to synthesize a person image with a new clothing, given the image of the person wearing a different clothing item and the new clothing item as inputs. Virtual try-on can be broken down into three main sub-tasks, namely image warping, image compositing, and synthesizing. The latter is very challenging as a synthetic image must preserve the person's identity, pose and shape. Also, the occluded body parts in a clothing item should be correctly synthesized. Moreover, the clothing image should accurately fit the pose and shape of a person, and the details of the clothing should also be preserved (i.e. logo, texture and embroidery). Prior work [8, 9, 11, 12, 13, 10, 12, 13, 14, 15] formulates virtual try-on as a supervised learning problem by following two major steps:

In order to address the aforementioned limitations, we introduce a self-supervised conditional generative adversarial network model, dubbed Fill In FAbrics (FIFA), which is a body-aware inpainting framework for image-based virtual try-on. The proposed FIFA framework can synthesize more realistic logo, texture and embroidery of the target clothing and also tackles well person images with complex poses (e.g., hands occluded). Our approach consists of a Fabricator and a unified virtual try-on pipeline with a Segmenter, Warper and Fuser. The Fabricator is used as a form of self-supervised pretraining for Warper. The goal of the Fabricator is to reconstruct full clothing details, given a partial input, enabling the model to learn the overall structure of the clothing (i.e. logo, texture, embroidery, full/short sleeves). To enforce global context at multiple scales for accurate modeling of the human and clothing interaction for compositing and synthesis, we also propose to use a multi-scale structural constraint to warp and refine the target clothing. The main contributions of this paper can be summarized as follows:

- We propose FIFA, a self-supervised conditional generative adversarial network model for virtual try-on, which can handle the complex pose of a reference person while preserving the target clothing details.
- We design a masked cloth modeling (MCM) objective to learn the overall structure of the clothing by predicting the full clothing image, given a masked input, for the downstream task of better target cloth warping and refinement.
- We show through experimental results and ablation studies that our model achieves competitive performance in comparison with strong baselines, yielding more realistic virtual try-on outputs.

2 Related Work

Image-Based Virtual Try-On. The basic objective of image-based virtual try-on is to synthesize a photo-realistic new image by overlaying a desired product image seamlessly onto the corresponding region of a clothed person. To achieve this goal, various image-based virtual try-on methods based on generative models have been proposed, Conditional Analogy Generative Adversarial Network (CA-GAN) [13], Virtual Try-On Network (VITON) [1],

Characteristic-Preserving Virtual Try-On (CP-VTON) network [22], CP-VTON+ [22], Disentangled Cycle-consistency Try-On Network (DCTON) [3], ClothFlow [11], SieveNet [12], Adaptive Content Generating and Preserving Network (ACGPN) [23], and Cloth Interactive Transformer (CIT) [22]. While these methods aim to handle complex textures on clothes and reduce artifacts in the final try-on results, they fail when the visual difference between the person image and target clothing is significant (e.g., changing long sleeve clothing items with short sleeve) and also tend to generate distorted arm regions. Furthermore, they fail to tackle person images with complex poses.

Masked Data Modeling. Masked data modeling has proven effective in natural language processing and computer vision [**b**, **13**, **19**, **25**]. Existing masked data modeling approaches include Context Encoders [**21**] and Masked Autoencoders (MAE) [**12**]. Our work differs from existing methods in two main aspects. First, we make predictions at a pixel level compared to predicting visual tokens [**1**]. Second, our encoder network is purely convolutional by design, and is not based on vision transformers, which have been shown to perform well only when pre-trained on large-scale image datasets such as the JFT-300M dataset [**b**].

3 Proposed Method

Problem Statement. Image-based virtual try-on aims to synthetically fit a target clothing onto a reference person while preserving photo-realistic details such as identity, pose and shape of the person, as well as texture and embroidery of the target clothing. More precisely, given a reference person image and a clothing image, the goal of our proposed FIFA model is to synthesize a new image of the same person wearing the target clothing such that the shape and pose of the person, as well as the details of the clothing are preserved.

3.1 Fill in Fabrics for Virtual Try-On

The proposed FIFA framework consists of a Fabricator and a unified pipeline consisting of a Segmenter, Warper and Fuser for virtual try-on, as shown in Figure 1. Given a partial input, we first use the Fabricator to reconstruct the full clothing details and learn the overall structure of the clothing (i.e. texture, full and half sleeve). This is used as a pretext task for the Warper. Second, we use Segmenter to predict the mask of the body parts of the reference person, as well as the masked target clothing regions. Third, we employ Warper to warp the target clothing image such that it fits the masked clothing region with the aim to capture the pose and shape of the reference person. Finally, Fuser integrates the outputs from Segmenter and Warper in order to synthesize the final try-on image.

Fabricator. The Fabricator aims to reconstruct (i.e. fill in fabrics) the full target clothing image $\hat{\mathbf{T}}_c$, given the partial target clothing $\mathbf{T}_{partial}$. To this end, the Fabricator learns to represent the overall structure of the clothing while reconstructing the missing regions (i.e. fill in correct pixels that make sense in the context). Inspired by the concept of image inpainting (i.e. the task of filling in holes in an image) using partial convolutions, where the convolution is masked and re-normalized to be conditioned on only valid pixels [**LN**], we construct $\mathbf{T}_{partial}$ from \mathbf{T}_c using masks of random streaks and holes of arbitrary shapes. In contrast to image inpainting, we formulate our objective as a masked cloth modeling problem, which can be regarded as a form of self-supervised pre-training for the downstream task of virtual try-on. More specifically, we train an encoder-decoder network \mathcal{F}_s to reconstruct the reconstructed



Figure 1: Schematic layout of the proposed FIFA framework for virtual try-on. Given a person image I and a clothing image T_c , FIFA synthesizes a try-on image I_t, where the person in image I is wearing the target clothing T_c . STN refers to the spatial transformer network, and \oplus denotes concatenation.

target clothing $\hat{\mathbf{T}}_c$, for a given $\mathbf{T}_{\text{partial}}$, with the goal to be close to the original target clothing image \mathbf{T}_c (i.e. non-masked clothing) by minimizing the L_1 error $\mathcal{E} = \|\hat{\mathbf{T}}_c - \mathbf{T}_c\|_1$.

Segmenter. The goal of the Segmenter is to preserve the body parts of the person during the synthesis process and also to accurately predict the semantic layout of the target clothing regions that are necessary for the Warper. Given a reference person image I and its associated mask M obtained via a publicly available human parser [12], the arms and torso regions are merged to form a fused map M_{fused} . A conditional generative adversarial network (CGAN) G_p is then trained to generate a different person body part mask M_{bp} , which is conditioned on M_{fused} , the 18-keypoint pose heatmap M_{pose} using out-of-the-box 2D pose estimator [2, 6], and the target clothing image T_c . To generate the target clothing region M_{cloth} , another CGAN G_c is trained by combining M_{bp} , M_{pose} and T_c . Hence, in the Segmenter there are two CGANs in which the discriminator is similar to pix2pixHD [22] and the generator is a Residual U-Net architecture [60] built on top of the U-Net model [23] with residual connections [10]. This not only helps retain fine-grained features and predict accurate body part masks, but also helps generate better try-on results. For a given CGAN (i.e. G_p or G_c), the adversarial loss is given by

$$\mathcal{L}_{\text{CGAN}} = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})}[\log D(\mathbf{x}|\mathbf{y})] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})}[\log(1 - D(G(\mathbf{z}|\mathbf{y})))], \quad (1)$$

where *G* and *D* are the generator and discriminator, **x** and **y** are the input and ground-truth mask, and **z** is a noise prior drawn from a standard normal distribution. A CGAN is a type of GAN that takes advantage of auxiliary information during the training process. To train a CGAN, we train the generator and discriminator simultaneously to maximize the performance of both. In simple terms, the goal of the generator is to generate data that the discriminator classifies as "real", whereas the objective of the discriminator is to not be "fooled" by the generator. In other words, the generator and discriminator follow the two-player min-max game with \mathcal{L}_{CGAN} as a function of *G* and *D*.

In order to enforce consistency at the pixel-level, we also use the pixel-wise cross-entropy loss \mathcal{L}_{CE} for better semantic segmentation results from the generator. Therefore, the overall objective is defined as

$$\mathcal{L}_{\text{mask}} = \alpha_1 \mathcal{L}_{\text{CGAN}} + \alpha_2 \mathcal{L}_{\text{CE}},\tag{2}$$

where α_1 and α_2 are nonnegative regularization parameters, which control the contribution of each loss term. Following previous work [23], we set α_1 and α_2 to 1 and 10, respectively, in our experiments.

We employ Warper to naturally deform the target clothing to fit the mask of the Warper. clothing region with respect to the pose of the person, as well as to preserve the texture and embroidery of the target clothing. While the Adaptive Content Generating and Preserving Network (ACGPN) [23] for virtual try-on has been shown effective at predicting the semantic layout of the reference image, it fails, however, to preserve complex poses, logo, texture and embroidery of the target clothing. This is largely due to the fact that ACGPN employs the Spatial Transformer Network (STN) [1] with Thin Plate Splines (TPS) [2] and an additional refinement network U-Net [23]. To address these limitations, we design a masked cloth modeling objective (MCM) when training the Warper to better preserve logo, texture and embroidery of the target clothing. More specifically, we transfer the learned representations in \mathcal{F}_s from Fabricator to the refinement network. We also incorporate a multi-scale structural constraint (MSC) to enforce global context at multiple scales for better warping of the target clothing according to the pose and shape of the person. Our strategy of training Warper yields better warped target clothes, which have fine details (i.e. logo, texture and embroidery), and is especially effective at handling complex poses.

Given the target clothing region \mathbf{M}_{cloth} and target clothing image \mathbf{T}_c , the goal of Warper is to deform \mathbf{T}_c such that it fits \mathbf{M}_{cloth} . STN first warps the clothing to \mathbf{T}_{warped} . This is further refined using \mathbf{T}_{warped} as input to the refinement network with the goal to generate more details (i.e. logo, texture, embroidery). In a similar vein to [27, 23], composition is then performed on the output of the refinement network with \mathbf{M}_{cloth} to output the final refined clothing $\mathbf{T}_{refined}$. The overall loss for the STN in Warper is an unweighted combination of the \mathcal{L}_{CGAN} loss and a second-order difference constraint [23]. The losses for the refinement network (i.e. pre-trained from the encoder-decoder network \mathcal{F}_s) are \mathcal{L}_{CGAN} and the perceptual \mathcal{L}_{VGG} loss [16]. This VGG perceptual loss helps ensure the target clothing and its warped version contain the same semantic content. In addition, we introduce a multi-scale structural constraint to enforce global context at multiple scales during training. Therefore, the overall loss function is defined as

$$\mathcal{L}_{\text{refined}} = \beta_1 \mathcal{L}_{\text{CGAN}} + \beta_2 \mathcal{L}_{\text{VGG}} + \beta_3 \mathcal{L}_{\text{MS-SSIM}}$$
(3)

where β_1 , β_2 and β_3 are regularization parameters, which are set to 0.2, 20 and 15, respectively, in our experiments. $\mathcal{L}_{MS-SSIM}$ is the multi-scale structural similarity constraint [51]. The Warper benefits from the MCM objective and is able to better preserve the logo, texture and embroidery of the target clothing. It also benefits from MSC to enforce global context in order to ensure better warping of the target clothing according to the pose and shape of the person. This in turn helps produce improved try-on results in Fuser.

Fuser. The Fuser merges the target clothing region, refined clothing image, a composited body part mask and a body part image with original clothing region masked out in order to produce the final try-on image. First, the Fuser generates a composited body part mask to remove or preserve the non-target body parts, which correspond, in most cases, to the arms of the person. This is then used in the second stage to determine which parts to preserve or generate when synthesizing the final try-on results. Given the original body part mask M_{obp} , the clothing mask M_{oc} from M (i.e. head, arms, torso removed), M_{bp} and M_{cloth} from Segmenter, the composited body part mask M_{comp} is given by

$$\mathbf{M}_{\text{comp}} = \left(\left(\mathbf{M}_{\text{bp}} \odot \mathbf{M}_{\text{oc}} \right) + \mathbf{M}_{\text{obp}} \right) \odot \left(\mathbf{J} - \mathbf{M}_{\text{cloth}} \right), \tag{4}$$

where \odot denotes element-wise multiplication and **J** is an all-ones matrix. As this step takes an input from Segmenter, it is crucial to produce accurate segmentation maps of M_{bp} and M_{cloth} for better compositing. We also perform compositing on **I** to get the body part image with I_{nc} being the original clothing region masked out as follows:

$$\mathbf{I}_{nc} = (\mathbf{I} - \mathbf{M}_{oc}) \odot (\mathbf{J} - \mathbf{M}_{cloth}).$$
(5)

Hence, given $\mathbf{T}_{refined}$ from Warper, \mathbf{M}_{cloth} from Segmenter, \mathbf{M}_{comp} and \mathbf{I}_{nc} , we train a CGAN G_m to predict the final try-on image \mathbf{I}_t by minimizing the following loss function

$$\mathcal{L}_{\text{fuser}} = \gamma_1 \mathcal{L}_{\text{CGAN}} + \gamma_2 \mathcal{L}_{\text{VGG}},\tag{6}$$

where the hyper-parameters γ_1 and γ_2 are set to 1 and 10, respectively, in our experiments.

4 **Experiments**

We conduct extensive experiments to assess the performance of the proposed FIFA framework in comparison with competing baseline models for virtual try-on. Experimental details and additional results and ablation studies are provided in the supplementary material. Code is available at: https://github.com/hasibzunair/fifa-tryon

4.1 Experimental Setup

Datasets. We demonstrate and analyze the performance of our model on two virtual try-on datasets: VITON and DecaWVTON.

- VITON: This dataset consists of 16,253 pairs of front-view women images and front-view top clothing images split into a training set of 14,221 pairs and a test set of 2,032 pairs. To evaluate the capability of virtual try-on methods in handling different poses of a person, we divide the VITON test set into three subsets of easy, medium and hard cases according to the human pose in the reference images. These test subsets are denoted as VITON-E, VITON-M and VITON-H for easy, medium and hard, respectively [23].
- **DecaWVTON:** To demonstrate the generalizability of FIFA to in-the-wild images, we use DecaWVTON, a proprietary dataset comprised of images with complex poses and clothing not present in the VITON dataset (e.g., turtle neck). Also, the clothing images are rotated, whereas VITON consists of only front-view clothing images. In many cases, the head portion is cut out (i.e. either fully or partially), whereas in VITON the person images consist of full faces.

Baselines. We evaluate the performance of our proposed virtual try-on model against recent state-of-the-art techniques, including CA-GAN [13], VITON [1], CP-VTON [26], CP-VTON+ [21], SieveNet [12], and segmentation based methods such as VTNFP [29] and ACGPN [23], as well as flow based methods such as ClothFlow [11]. We also compare our model against a cycle-consistency based approach DCTON [2] and a transformer based method CIT [22].

Evaluation Metrics. Following previous work [21, 23], we use the Structural SIMilarity (SSIM) that captures image level similarity and the Frechet Inception Distance (FID) that

captures the distributional similarity. Both metrics are commonly used for benchmarking virtual try-on methods to quantify the visual difference between the generated and real reference images. Higher scores of SSIM and lower scores of FID indicate higher quality of the synthesized results. It is important to mention that while computing the SSIM and FID metrics, the target clothing items are the same as in the reference person as it is not possible to acquire ground truth images for try-on results.

Implementation Details. All experiments are performed on a Linux workstation running 4.8Hz, 64GB RAM and a single NVIDIA RTX 3080 GPU. Experiments are conducted using Python programming language and PyTorch deep learning framework. A full training of FIFA, along with the Fabricator on the VITON dataset, takes roughly seven days. During training, the target clothing item is the same as the one in the reference person image, as it is not possible to acquire triplets to compute the loss with respect to the ground truth.

4.2 Qualitative Results

In Figure 2, we visually compare the performance of our proposed model with CP-VTON+ and ACGPN, which are state-of-the-art virtual try-on baselines. Each row shows a person virtually trying on different clothing items. As can be seen in the first row of Figure 2, when the pose of the reference person is complex (i.e. standing with arms behind the body), the baseline models either remove body regions, fail to warp short sleeve shirt, or add unrealistic body parts. These baselines are also unable to capture the global structure and semantics, which are needed for warping short sleeve shirts when the reference person is wearing a long sleeve shirt. This is due, in large part, to the limited capability of the warping strategies used in these baselines. The second and third rows of Figure 2 show cases where the target clothing items are of complex texture (i.e. printed patterns, long sleeve, and shirt with logo) and embroidery (i.e. stripes). In these cases, CP-VTON+ fails to distinguish between the front and back part of the clothing regions, does not preserve the logo of the target clothes, and yields blurry results at the clothing and person body boundaries. While ACGPN produces non-blurry results, it fails to preserve the complex embroidery of the target clothing, and does not accurately warp long sleeve target clothing items, resulting in incomplete sleeves. In the last row of Figure 2, we can observe artifacts and mix-up of front and back part of the clothing in the images generated by the baseline methods. Also, both CP-VTON+ and ACGPN fail to capture the v-shaped structure of the target clothing, and do not accurately warp tank tops with very thin straps, resulting in either blurry or distorted clothing structure. Overall, these baselines fail to preserve the complex pose of the reference person, the complex texture and embroidery of the target clothing, and also the complex clothing types.

By comparison, our FIFA method is able to warp the target clothing in the case of complex poses, and preserves well the body parts. It benefits from the synergy between the MCM objective and the MSC constraint, which help preserve the pose of a person, capture the fine details of the target clothing (i.e. logo and embroideries), as well as the global structure of the clothing (i.e. front and back part of clothing, v-shaped collar). Moreover, FIFA benefits from residual blocks (RBs) to better predict the semantic layout of the body parts, resulting in realistic try-on results. In summary, this helps not only preserve the logo, texture, embroidery and the type of target clothing, but also yields an output having less artifacts and retains clear body parts, achieving more realistic try-on results. We also find that FIFA is able to better preserves the skin color of the person and accurately synthesizes the person's body parts, which were initially occluded. In addition, it can distinguish between the front



Figure 2: Given a pair of a reference person image and a target clothing image, our FIFA model successfully synthesizes virtual try-on images. Compared to the baselines, FIFA is able to better handle complex poses and also retains photo-realistic details such as logo, texture, embroidery and structure (e.g., collar shape) of the target clothing.

and back part of clothing items.

It is worth pointing out that some examples in Figure 2 seem to have color mismatch between synthetic clothes and target. We hypothesize that this might be attributed to the Warper, which in some cases produces blurry target clothing outputs. While using the multi-scale structural constraint (MSC) in Warper can output fine details of clothing, we argue that designing perceptually motivated loss functions may further improve the results.

4.3 Quantitative Results

Table 1 shows that FIFA consistently outperforms all baselines, achieving relative improvements of 4.85%, 4.22%, 4.64% and 4.47% over the strongest ACGPN baseline on all (VI-TON), easy, medium and hard cases in terms of the SSIM metric. FIFA also outperforms ACGPN with a substantial relative improvement of 19.11% in terms of FID.

Interestingly, our FIFA model yields significant relative improvements of 5.10% and 43.16% over ClothFlow in terms of SSIM and FID, respectively. It is worth pointing out that ClothFlow operates on streams (i.e. optimal flow maps) to predict the movement of clothes and is computationally expensive, while our model is a purely image-based virtual try-on approach operating on image pixels. Our method also outperforms the transformer based CIT baseline with a relative improvement of 7.13% in terms of SSIM. This better performance of our approach is significant because transformers are built on self-attention operations and are quite strong in modeling the global context between the person and target clothing. In addition, transformers in computer vision tasks perform well only when pre-trained on a large cohort of images such as the JFT-300M dataset, which is comprised of 18K classes and 303M high-resolution images [**B**].

4.4 Ablation Study

Effectiveness of Masked Cloth Modeling (MCM). Figure 3 illustrates the benefit of using the MCM objective in preserving the pose and logo, as well as in accurately warping the target clothing. As can be seen, MCM is able to preserve the logo of the target clothing,

SSIM (↑)					_
Method	VITON	VITON-E	VITON-M	VITON-H	FID (\downarrow)
CA-GAN [0.740	-	-	-	47.34
VITON 🖪	0.783	0.787	0.779	0.779	55.71
CP-VTON [26]	0.745	0.753	0.742	0.729	24.43
VTNFP [29]	0.803	0.810	0.801	0.788	-
ClothFlow [0.843	-	-	-	23.68
CP-VTON+ [20]	0.750	-	-	-	21.08
SieveNet [💶]	0.837	-	-	-	26.67
ACGPN [28]	0.845	<u>0.854</u>	0.841	0.828	16.64
DCTON [8]	0.830	-	-	-	14.82
CIT [Z]	0.827	-	-	-	-
FIFA (Ours)	0.886	0.890	0.880	0.865	13.46

Table 1: Performance comparison of FIFA and state-of-the-art methods on the VITON, VITON-E, VITON-M and VITON-H test sets using SSIM and FID scores. FIFA consistently outperforms the baselines across easy, medium and hard cases. Boldface numbers indicate the best performance, whereas the best baselines are underlined.

whereas without MCM the logo is completely lost. MCM also helps in accurately preserving or synthesizing body parts in complex poses (i.e. body aware), as well as in accurately warping the target clothing. Notice that without MCM, there is a problem of unnecessarily editing regions of the target clothing (i.e. making half sleeve shirt a full sleeve). This is largely attributed to the richer learning signal provided by MCM rather than just using the supervised objective of predicting the warped clothing, which fits the reference person, enabling our approach to accurately model the interactions between the target clothing and reference person clothing.



Figure 3: Warped target clothing results, demonstrating the effectiveness of the MCM objective in Warper. Warper with MCM is capable of handling complex poses (i.e. body-aware) and preserving the logo and embroidery of the clothing.

 to capture the global context and semantics, which are needed for preserving the shape of a person and also for realistically synthesizing body parts. The per-pixel-based loss function \mathcal{L}_1 measures the distance between pixels and does not enforce any global constraint. On the other hand, the perceptual loss \mathcal{L}_{VGG} quantifies the similarity between the reconstructed and ground-truth images, but only at a latent representation level (i.e. computes the distance of the features extracted by VGG-19 [23]). Also, it tends to generate artifacts [16], which is in line with our findings. We show that by adding MSC, our model is able to better tackle these issues and learns to exploit context at different scales, while the CP-VTON+ and ACGPN baselines introduce artifacts and do not preserve well the shape of the person.



Figure 4: Try-on results, demonstrating the effectiveness of MSC in Warper. Warper with MSC helps capture global context of the target clothing and preserves the shape of the person.

4.5 Generalization to In-The-Wild Virtual Try-On

To test the generalizability of virtual try-on models to in-the-wild images, we set up a challenging task where the results would better reflect the robustness on unseen data. We compare FIFA against the state-of-the-art ACGPN model [23] by training both methods on VI-TON and testing them on DecaWVTON. Results presented in the supplementary material demonstrate that FIFA yields substantial improvements over ACGPN in terms of SSIM and FID, indicating that FIFA is more robust to in-the-wild images for virtual try-on.

5 Conclusion

We introduced a body-aware self-supervised inpainting framework for image-based virtual try-on with a focus on tackling complex poses, learning the overall structure of clothing and incorporating global context. Our proposed FIFA model achieves significant improvements in the synthesized try-on image by not only retaining the logo, texture and embroidery of the clothing, but also able to better handle the complex poses, indicating that it is body aware, a crucial feature for photo-realistic virtual try-on. By combining the strengths of mask cloth modeling, multi-scale structural constraint and residual blocks, FIFA outperforms strong baselines on the VITON dataset across all, easy, medium and hard cases. In addition, we set up an evaluation framework for testing robustness of virtual try-on models to in-the-wild images and found that FIFA outperforms previous state-of-the-art methods by a significant margin.

References

- [1] Hangbo Bao, Li Dong, and Furu Wei. BEiT: BERT pre-training of image transformers. In *International Conference on Learning Representations*, 2022.
- [2] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. OpenPose: Realtime multi-person 2D pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2D pose estimation using part affinity fields. In Proc. IEEE Conference on Computer Vision and Pattern Recognition, 2017.
- [4] Seunghwan Choi, Sunghyun Park, Minsoo Lee, and Jaegul Choo. VITON-HD: Highresolution virtual try-on via misalignment-aware normalization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 14131–14140, 2021.
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pretraining of deep bidirectional transformers for language understanding. In Proc. Annual Conference of the North American Chapter of the Association for Computational Linguistics, pages 4171–4186, 2019.
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations*, 2021.
- [7] Jean Duchon. Splines minimizing rotation-invariant semi-norms in Sobolev spaces. In Constructive Theory of Functions of Several Variables, pages 85–100. Springer, 1977.
- [8] Chongjian Ge, Yibing Song, Yuying Ge, Han Yang, Wei Liu, and Ping Luo. Disentangled cycle consistency for highly-realistic virtual try-on. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 16928–16937, 2021.
- [9] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S Davis. VITON: An imagebased virtual try-on network. In Proc. IEEE Conference on Computer Vision and Pattern Recognition, pages 7543–7552, 2018.
- [10] Xintong Han, Xiaojun Hu, Weilin Huang, and Matthew R Scott. ClothFlow: A flowbased model for clothed person generation. In *Proc. IEEE International Conference on Computer Vision*, pages 10471–10480, 2019.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proc. IEEE Conference on Computer Vision and Pattern Recognition, pages 770–778, 2016.
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. arXiv preprint arXiv:2111.06377, 2021.
- [13] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In Advances in Neural Information Processing Systems, 2015.

- [14] Surgan Jandial, Ayush Chopra, Kumar Ayush, Mayur Hemani, Balaji Krishnamurthy, and Abhijeet Halwai. SieveNet: A unified framework for robust image-based virtual try-on. In *Proc. IEEE Winter Conference on Applications of Computer Vision*, pages 2182–2190, 2020.
- [15] Nikolay Jetchev and Urs Bergmann. The conditional analogy GAN: Swapping fashion articles on people images. In *Proc. IEEE International Conference on Computer Vision Workshops*, pages 2287–2292, 2017.
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Proc. European Conference on Computer Vision*, pages 694–711. Springer, 2016.
- [17] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Self-correction for human parsing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [18] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proc. European Conference on Computer Vision*, pages 85–100, 2018.
- [19] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781, 2013.
- [20] 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 Proc. IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2020.
- [21] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 2536–2544, 2016.
- [22] Bin Ren, Hao Tang, Fanyang Meng, Runwei Ding, Ling Shao, Philip HS Torr, and Nicu Sebe. Cloth interactive transformer for virtual try-on. arXiv preprint arXiv:2104.05519, 2021.
- [23] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [24] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for largescale image recognition. In *International Conference on Learning Representations*, 2015.
- [25] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In Proc. IEEE Winter Conference on Applications of Computer Vision, pages 2149–2159, 2022.
- [26] Bochao Wang, Huabin Zheng, Xiaodan Liang, Yimin Chen, Liang Lin, and Meng Yang. Toward characteristic-preserving image-based virtual try-on network. In *Proc. European Conference on Computer Vision*, pages 589–604, 2018.

- [27] 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 *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 8798–8807, 2018.
- [28] 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 *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 7850–7859, 2020.
- [29] Ruiyun Yu, Xiaoqi Wang, and Xiaohui Xie. VTNFP: An image-based virtual try-on network with body and clothing feature preservation. In *Proc. IEEE International Conference on Computer Vision*, pages 10511–10520, 2019.
- [30] Zhengxin Zhang, Qingjie Liu, and Yunhong Wang. Road extraction by deep residual U-Net. *IEEE Geoscience and Remote Sensing Letters*, 15(5):749–753, 2018.
- [31] Hang Zhao, Orazio Gallo, Iuri Frosio, and Jan Kautz. Loss functions for image restoration with neural networks. *IEEE Transactions on Computational Imaging*, 3(1):47–57, 2016.