

# Do Not Mask What You Do Not Need to Mask: a Parser-Free Virtual Try-On

Thibaut Issenhuth<sup>1</sup>, Jérémie Mary<sup>1</sup>, and Clément Calauzènes<sup>1</sup>

Criteo AI Lab, Paris, France

{t.issenhuth,j.mary,c.calauzenes}@criteo.com

**Abstract.** The 2D virtual try-on task has recently attracted a great interest from the research community, for its direct potential applications in online shopping as well as for its inherent and non-addressed scientific challenges. This task requires fitting an in-shop cloth image on the image of a person, which is highly challenging because it involves cloth warping, image compositing, and synthesizing. Casting virtual try-on into a supervised task faces a difficulty: available datasets are composed of pairs of pictures (cloth, person wearing the cloth). Thus, we have no access to ground-truth when the cloth on the person changes. State-of-the-art models solve this by masking the cloth information on the person with both a human parser and a pose estimator. Then, image synthesis modules are trained to reconstruct the person image from the masked person image and the cloth image. This procedure has several caveats: firstly, human parsers are prone to errors; secondly, it is a costly pre-processing step, which also has to be applied at inference time; finally, it makes the task harder than it is since the mask covers information that should be kept such as hands or accessories. In this paper, we propose a novel student-teacher paradigm where the teacher is trained in the standard way (reconstruction) before guiding the student to focus on the initial task (changing the cloth). The student additionally learns from an adversarial loss, which pushes it to follow the distribution of the real images. Consequently, the student exploits information that is masked to the teacher. A student trained without the adversarial loss would not use this information. Also, getting rid of both human parser and pose estimator at inference time allows obtaining a real-time virtual try-on.

**Keywords:** Virtual try-on. Teacher-student. Model distillation.

## 1 Introduction

A photo-realistic virtual try-on system would provide a significant improvement for online shopping. Whether used to create catalogs of new products or to propose an immersive environment for shoppers, it could impact e-commerce and open the door for automated image-editing possibilities.

Earlier work addresses this challenge using 3D measurements and model-based methods [9,11,28]. However, these are, by nature, computationally intensive and require expensive material, which would not be acceptable at scale



**Fig. 1:** Typical failure cases of the human parser. On the two first rows, it does not segment the person properly. On the third row, it masks the handbag which we would like to preserve in a virtual try-on. CP-VTON and our T-WUTON, which rely on the parsing information, are not robust to a bad parsing. However, the student model S-WUTON which is distilled from the human parser, pose estimator and T-WUTON, can preserve the person’s attributes and does not rely on the parsing information.

for shops. Recent works aim to leverage deep generative models to tackle the virtual try-on problem [6,13,19,38]. CAGAN [19] is a U-net based Cycle-GAN [17] approach. However, this method fails to generate realistic results since such networks cannot handle large spatial deformations. In VITON [13], the authors recast the virtual try-on as a supervised task. They propose to use a human parser and a pose estimator to mask the cloth in the person image and construct an agnostic person representation  $p^*$ . The human parser allows segmenting the upper-body and the cloth, while the pose estimator locates the keypoints (*i.e.* shoulders, wrists, etc.) of the person. Then, with  $p^*$  and the image of the original cloth  $c$  on a white background, they train a model in a fully supervised fashion to reconstruct  $p$ . Namely, they propose a coarse-to-fine synthesis strategy with shape context matching algorithm [3] to warp the cloth on the target person. To improve this model, CP-VTON [38] incorporates a convolutional geometric matcher [31], which learns geometric deformations (*i.e.* thin-plate spline transform [4]) that align the cloth with the person. State-of-the-art models are based on the supervised formulation of the virtual try-on task, which has some drawbacks. Human parsers and pose estimators are trained on other datasets and thus

fail in some situations (see Fig. 1, two first rows). Retraining them on fashion datasets would require similar labels of semantic segmentation or unsupervised domain adaptation methods. Even though they would still be imperfect. Moreover, for a virtual try-on, one wants to preserve person’s attributes like handbags or jewels. When constructing  $p^*$ , these person’s attributes are masked and can not be preserved, such as the partially masked handbag on the third row of Fig. 1. Finally, the human parsing and pose estimation are the wall clock bottleneck of the pipeline.

In our work, we distill [15] the standard pipeline of virtual try-on composed of human parser, pose estimator, and synthesis modules in the synthesis modules. Namely, we train a student synthesizer with the outputs of a pre-trained standard virtual try-on pipeline. To force the student to use information that is masked to the teacher, we also train the student with an adversarial loss. The distillation process allows us to remove the need for human parsing and pose estimation at inference time, which improves image quality and speeds up the computations from 6FPS to 77FPS. In Fig. 1, we show visual results of a baseline CP-VTON, our teacher model T-WUTON and our student model S-WUTON. Since S-WUTON does not rely on human parsing, it is robust to parsing errors and preserves a person’s attributes such as fingers or handbags.

Additionally, to build an efficient teacher model, we propose an improved architecture for virtual try-on, a Warping U-Net for a Virtual Try-On (WUTON). Our architecture is composed of two modules: a convolutional geometric matcher [31] and a U-net generator with a siamese encoder, where the former warps the feature maps of the latter. The architecture is trained end-to-end, which leads to high-quality synthesized images.

We demonstrate the benefit of our method with several experiments on a virtual try-on dataset, with quantitative and visual results, and a user study.

## 2 Problem statement and related work

Given the 2D images  $p \in \mathbb{R}^{h \times w \times 3}$  of a person and  $c \in \mathbb{R}^{h \times w \times 3}$  of a clothing item, we want to generate the image  $\tilde{p} \in \mathbb{R}^{h \times w \times 3}$  where a person  $p$  wears the cloth  $c$ . The task can be separated in two parts : the geometric deformation  $T$  required to align  $c$  with  $p$ , and the refinement that fits the aligned cloth  $\tilde{c} = T(c)$  on  $p$ . These two sub-tasks can be modelled with learnable neural networks, *i.e.* spatial transformers networks  $STN$  [18,31] that output parameters  $\theta = STN(p, c)$  of geometric deformations, and conditional generative networks  $G$  that give  $\tilde{p} = G(p, c, \theta)$ .

Because it would be costly to construct a dataset with  $\{(p, c), \tilde{p}\}$  triplets, previous works [13,38] propose to use an agnostic person representation  $p^* \in \mathbb{R}^{h \times w \times c}$  where the clothing items in  $p$  are hidden but identity and shape of the persons are preserved.  $p^*$  is built with pre-trained human parsers and pose estimators :  $p^* = h(p)$ . These triplets  $\{(p^*, c), p\}$  allow to train for reconstruction.  $(p^*, c)$  are the inputs,  $\tilde{p}$  the output and  $p$  the ground-truth. We finally have the

conditional generative process :

$$\tilde{p} = G(\underbrace{h(p)}_{\text{agnostic person}}, \underbrace{c}_{\text{cloth}}, \underbrace{STN(h(p), c)}_{\text{geometric transform}}) \quad (1)$$

Although it eases the training of  $G$ ,  $h$  is a bottleneck in the virtual try-on pipeline. We will show that we can train a student model with synthetic triplets  $\{(p, c), \tilde{p}\}$ , where  $\tilde{p}$  comes from our pre-trained teacher generative model in Eq. 1. This allows to remove the need for  $h$  at inference time for the student model:

$$\hat{p} = G_s(\underbrace{p}_{\text{original person}}, \underbrace{c}_{\text{cloth}}, \underbrace{STN_s(p, c)}_{\text{geometric transform}}) \quad (2)$$

where  $G_s$  and  $STN_s$  are the student modules and  $\hat{p}$  the generated image.

**Conditional image generation.** Generative models for image synthesis have shown impressive results with adversarial training [8]. Combined with deep networks [29], this approach has been extended to conditional image generation in [27] and performs increasingly well on a wide range of tasks, from image-to-image translation [17,45] to video editing [34]. However, as noted in [26], these models cannot handle large spatial deformations and fail to modify the shape of objects, which is necessary for a virtual try-on.

**Appearance transfer.** Close to the virtual try-on task, some research focus on human appearance transfer. Given two images of different persons, the goal is to transfer the appearance of a part of the person A on the person B. Approaches using pose and appearance disentanglement [24,25] fit this task but others are specifically designed for it. SwapNet [30] is a dual path network which generates a new human parsing of the reference person and region of interest pooling to transfer the texture. In [40], the method relies on DensePose information [1], which provides a 3D surface estimation of a human body, to perform a warping and align the two persons. The transfer is then done with segmentation masks and refinement networks. However, the warping relies on matching source and target pose, which is not feasible for the virtual try-on task.

**Virtual try-on.** Most of the approaches for a virtual try-on system come from computer graphics and rely on 3D measurements or representations. Drape [9] learns a deformation model to render clothes on 3D bodies of different shapes. In [11], Hahn et al. use subspace methods to accelerate physics-based simulations and generate realistic wrinkles. ClothCap [28] aligns a 3D cloth-template to each frame of a sequence of 3D scans of a person in motion. However, the use of 3D scans is expensive and thus not doable for online users.

The task we are interested in is the one introduced in CAGAN [19] and further studied by VITON [13] and CP-VTON [38], which we defined in the problem statement. In CAGAN [19], Jetchev et al. propose a cycle-GAN approach that requires three images as input: the reference person, the cloth worn by the person

and the target in-shop cloth. Thus, it limits its practical uses. To facilitate the task, VITON [13] introduces the supervised formulation of the virtual try-on, as described above. Their pipeline separates the task in sub-tasks: constructing the agnostic person representation (*i.e.* mask the area to replace but preserve body shape), warping the cloth and compositing the final image. Based on the agnostic person representation  $p^*$  and the cloth image  $c$ , the VITON model performs a generative composition between the warped cloth and a coarse result. The warping is done with a non-parametric geometric transform [3]. To improve this model, CP-VTON [38] incorporates a learnable geometric matcher *STN* [31]. The *STN* is trained to align  $c$  on  $p$  with a L1 loss on paired images. However, the L1 loss is overwhelmed with the white background and the solid color parts of clothes. Thus, it faces difficulties to align patterns and to preserve inner structure of the cloth. In VTNFP [42] and ClothFlow [12], a module generating the new human parsing is added. It allows to better preserve body parts and edges, but at an increased computational cost. Moreover, ClothFlow [12] replaces the TPS warping by a dense flow from the target cloth to the person. All these recent works [13,38,12,42] rely on pre-trained human parser and pose estimator.

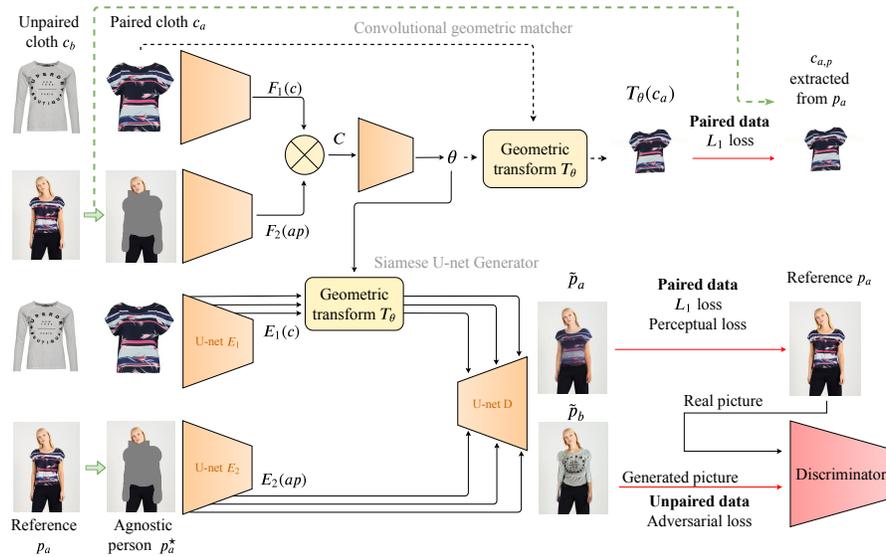
Recent work MG-VTON [6] extends the task to a multi-pose virtual try-on system, where they also change the pose of the reference person. Similarly to [5,42,6], they add a module generating the new human parsing, based on input and target pose information.

### 3 Our approach

Our task is to build a virtual try-on system that is able to fit a given in-shop cloth on a reference person. In this work, we build a virtual try-on that does not rely on a human parser nor a pose estimator for inference. To do so, we use a teacher-student approach to distill the standard virtual try-on pipeline composed of human parser, pose estimator, and synthesis module in the synthesis module.

In Section 3.1, we detail the architecture of our synthesis module WUTON. It is trainable end-to-end and composed of two existing modules: a convolutional geometric matcher *STN* [31] and a U-net [32] with siamese encoder whose skip connections from the cloth encoder to the decoder are deformed by *STN*. We then explain its training procedure in the standard supervised setting, which gives the teacher T-WUTON.

We finally explain our distillation process. Once the first generative model is trained, the pipeline  $\{h, \text{T-WUTON}\}$  becomes a teacher model for a student model S-WUTON by constructing synthetic triplets  $\{(p, c), \tilde{p}\}$ . These serve to supervise the training of S-WUTON, which hence does not need a human parser to pre-process the image and construct the agnostic person representation. Importantly, S-WUTON also learns from an adversarial loss so it does not only follow the teacher’s distribution and it can learn to preserve a person’s attributes.



**Fig. 2:** The teacher T-WUTON : our proposed end-to-end warping U-net architecture. Dotted arrows correspond to the forward pass only performed during training. Green arrows are the human parser, red ones are the loss functions. The geometric transforms share the same parameters but do not operate on the same spaces. The different training procedure for paired and unpaired pictures is explained in Section 3.2.

### 3.1 WUTON architecture

Our warping U-net is composed of two connected modules, as shown in Fig. 2. The first one is a convolutional geometric matcher, which has a similar architecture as [31,38]. It outputs the parameters  $\theta$  of a geometric transformation, a TPS transform in our case. This geometric transformation aligns the in-shop cloth image with the reference person. However, in contrast to previous work [6,13,38], we use the geometric transformation on the feature maps of the generator rather than at a pixel-level. Thus, we learn to deform the feature maps that pass through the skip connections of the second module, a U-net [32] generator which synthesizes the output image  $\tilde{p}$ .

The architecture of the convolutional geometric matcher is taken from CP-VTON [38], which reuses the generic geometric matcher from [31]. It is composed of two feature extractors  $F_1$  and  $F_2$ , which are standard convolutional neural networks. The local vectors of feature maps  $F_1(c)$  and  $F_2(p^*)$  are then L2-normalized and a correlation map  $C$  is computed as follows:

$$C_{ijk} = F_{1i,j}(c) \cdot F_{2m,n}(p^*) \quad (3)$$

where  $k$  is the index for the position  $(m, n)$ . This correlation map captures dependencies between distant locations of the two feature maps, which is useful to align the two images.  $C$  is the input of a regression network, which outputs the parameters  $\theta$  and allows to perform the geometric transformation  $T_\theta$ . We use

TPS transformations [4], which generate smooth sampling grids given control points. Each scale of the U-net is transformed with the same parameters  $\theta$ .

The input of the U-net generator is also the tuple of pictures  $(p^*, c)$ . Since these two images are not spatially aligned, we cannot simply concatenate them and feed a standard U-net. To alleviate this, we use two different encoders  $E_1$  and  $E_2$  processing each image independently and with non-shared parameters. Then, the feature maps of the in-shop cloth  $E_1(c)$  are transformed at each scale  $i$ :  $E_1^i(c) = T_\theta(E_1^i(c))$ . Then, the feature maps of the two encoders are concatenated and feed the decoder at each scale. With aligned feature maps, the generator is able to compose them and to produce realistic results. Feature maps warping was also proposed in [5,35]. We use instance normalization in the U-net generator, which is more effective than batch normalization [16] for image generation [37].

### 3.2 Training procedure of the teacher model

We will now detail the training procedure of T-WUTON, *i.e.* the data representation and the different loss functions of the teacher model.

While previous works use a rich person representation with more than 20 channels representing human pose, body shape and the RGB image of the head, we only mask the upper-body of the reference person. Our agnostic person representation  $p^*$  is thus a 3-channel RGB image with a masked area. We compute the upper-body mask from pose and body parsing information provided by a pre-trained neural network from [23]. Precisely, we mask the areas corresponding to the arms, the upper-body cloth and a bounding box around the neck.

Using the dataset from [6], we have pairs of in-shop cloth image  $c_a$  and a person wearing the same cloth  $p_a$ . Using a human parser and a human pose estimator, we generate  $p_a^*$ . From the parsing information, we can also isolate the cloth on the image  $p_a$  and get  $c_{a,p}$ , the cloth worn by the reference person. Moreover, we get the image of another in-shop cloth  $c_b$ . The inputs of our network are the two tuples  $(p_a^*, c_a)$  and  $(p_a^*, c_b)$ . The outputs are respectively  $(\tilde{p}_a, \theta_a)$  and  $(\tilde{p}_b, \theta_b)$ .

The cloth worn by the person  $c_{a,p}$  allows us to guide directly the geometric matcher with a  $L_1$  loss:

$$L_{warp} = \|T_{\theta_a}(c_a) - c_{a,p}\|_1 \quad (4)$$

The image  $p_a$  of the reference person provides a supervision for the whole pipeline. Similarly to CP-VTON [38], we use two different losses to guide the generation of the final image  $\tilde{p}_a$ , the pixel-level  $L_1$  loss  $\|\tilde{p}_a - p_a\|_1$  and the perceptual loss [20]. We focus on  $L_1$  losses since they are known to generate less blur than  $L_2$  for image generation [44]. The latter consists of using the features extracted with a pre-trained neural network, VGG [36] in our case. Specifically, our perceptual loss is:

$$L_{perceptual} = \sum_{i=1}^5 \|\phi_i(\tilde{p}_a) - \phi_i(p_a)\|_1 \quad (5)$$

where  $\phi_i(I)$  are the feature maps of an image  $I$  extracted at the  $i$ -th layer of the VGG network. Furthermore, we exploit adversarial training to train the network to fit  $c_b$  on the same agnostic person representation  $p_a^*$ , which is extracted from a person wearing  $c_a$ . This is only feasible with an adversarial loss, since there is no available ground-truth for this pair  $(p_a^*, c_b)$ . Thus, we feed the discriminator with the synthesized image  $\tilde{p}_b$  and real images of persons from the dataset. This adversarial loss is also back-propagated to the convolutional geometric matcher, which allows to generate much more realistic spatial transformations. We use the relativistic adversarial loss [21] with gradient-penalty [2,10], which trains the discriminator to predict relative realness of real images compared to synthesized ones. Finally, we optimize with Adam [22] the following objective function:

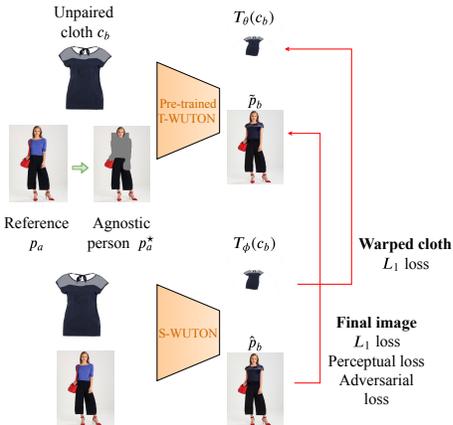
$$L = \lambda_w L_{warp} + \lambda_p L_{perceptual} + \lambda_{L_1} L_1 + \lambda_{adv} L_{adv} \quad (6)$$

### 3.3 Training procedure of the student model

We propose to use a teacher-student approach to distill the pipeline composed of  $\{h, \text{T-WUTON}\}$  in a single student WUTON (S-WUTON). Indeed, our pre-trained T-WUTON is able to generate realistic images and geometric deformations of clothes on images pre-processed by  $h$ . We leverage it and use it as a way to construct generated triplets  $\{(p_a, c_b), \tilde{p}_b\}$ , where  $\tilde{p}_b$  is the image synthesized by T-WUTON. With this pre-trained model, we can supervise the training of a student model S-WUTON. This allows to train the student model on the initial task of changing the cloth rather than reconstructing the upper-body. The student model has the exact same architecture than T-WUTON but different inputs and ground-truth. Hence, its inputs are  $(p_a, c_b)$ , where  $p_a$  is the non-masked image of a person. Having this non-masked image as input, the student model does not need a human parser for pre-processing images. The ground-truth of S-WUTON are the outputs of T-WUTON, for both the warped cloth  $T_\theta(c_b)$  and the final synthesized image  $\tilde{p}_b$ . The training scheme of the student model S-WUTON is shown in Fig. 3.

More precisely, let us define the inputs-outputs of the teacher and student model:  $(\tilde{p}_b, \phi) = \text{S-WUTON}(p_a, c_b)$  and  $(\tilde{p}_b, \theta) = \text{T-WUTON}(h(p_a), c_b)$ . Then, the loss functions of S-WUTON are:

$$L_{warp} = \|T_\phi(c_b) - T_\theta(c_b)\|_1 \quad (7)$$



**Fig. 3:** S-WUTON: our training scheme allowing to remove the need for a human parser at inference time. We use human parser and pre-trained T-WUTON to generate synthetic ground-truth for a student model S-WUTON.

$$L_{perceptual} = \sum_{i=1}^5 \|\phi_i(\hat{p}_b) - \phi_i(\tilde{p}_b)\|_1 \quad (8)$$

$$L_1 = \|\hat{p}_b - \tilde{p}_b\|_1 \quad (9)$$

Finally, the total loss of the student model is:

$$L = \lambda_w L_{warp} + \lambda_p L_{perceptual} + \lambda_{L_1} L_1 + \lambda_{adv} L_{adv} \quad (10)$$

The adversarial loss  $L_{adv}$  is independent from T-WUTON. Here, we also use the relativistic loss with gradient penalty on the discriminator. The real data consists of images of persons from the dataset  $p_a$ , and the fake data corresponds to the synthesized images  $\hat{p}_b$ . Notice that without the adversarial loss, it would be a standard teacher-student setting, where the student is only guided by the outputs of the teacher. In our case, the discriminator (*i.e.*  $L_{adv}$ ) helps S-WUTON to be close to the real data distribution, and not only to the teacher’s distribution. As shown by the ablation study in Section 4.6, it is an important component and is necessary to make S-WUTON exploit the components that are masked from T-WUTON (*e.g.* hands).

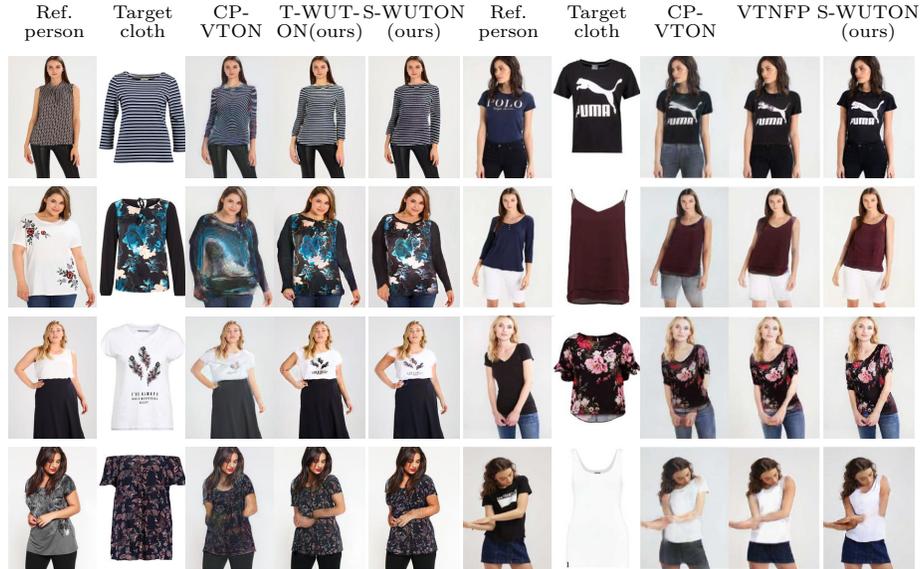
## 4 Experiments and analysis

We first describe the dataset. We then compare our approach with CP-VTON [38], a current state-of-the-art for the virtual try-on task. We present visual and quantitative results proving that S-WUTON achieves state-of-the-art results, and that the distillation process allows to improve image quality. We show that this stands for several metrics, and with a user study. We then provide a comparison of the runtime of virtual try-on algorithms on a Tesla NVIDIA V100 GPU. The teacher-student distillation allows to decrease the runtime by an order of magnitude. Finally, we outline the importance of the adversarial loss in our teacher-student setting.

We also show some visual comparisons with recent work VTNFP [42]. Images are taken from their paper. However, since their model is not available, we could not compute the other metrics. We provide more visual comparisons with VTNFP and ClothFlow [12] in supplementary material.

### 4.1 Dataset

For copyright issues, we can not use the dataset from VITON [13] and CP-VTON [38]. Instead, we leverage the *Image-based Multi-pose Virtual try-on* dataset. This dataset contains 35,687/13,524 person/cloth images at 256x192 resolution. 4175 pairs are kept for test so the cloth was not seen during training. A random shuffle of these pairs produces the unpaired person/cloth images. For each in-shop cloth image, there are multiple images of a model wearing the given cloth from different views and in different poses. We remove images tagged as back images since the



**Fig. 4:** On the left side, comparison of our method with CP-VTON [38]. For fairness, the two methods are trained on the same dataset and on the same agnostic person representation  $p^*$ . More examples are provided in supplementary material. On the right side, comparison with recent work VTNFP. Except for S-WUTON’s column, images are taken from their paper.

in-shop cloth image is only from the front. We process the images with a neural human parser and pose estimator, specifically the joint body parsing and pose estimation network [7,23].

## 4.2 Visual results

Visual results of our method and CP-VTON are shown in Fig. 4. On the left side, images are computed from models trained on MG-VTON dataset, with  $p_{t-wuton}^*$  representation for T-WUTON and CP-VTON for fairness. On the right side, images are taken from VTNFP paper [42]. There, CP-VTON and VTNFP were trained on the original dataset from VITON, and CP-VTON uses  $p_{cp-vton}^*$ . More images from S-WUTON are provided in Fig. 1 and Fig. 6.

CP-VTON has trouble to realistically deform and render complex patterns like stripes or flowers. Control points of the  $T_\theta$  transform are visible and lead to unrealistic curves and deformations on the clothes. Also, the edges of cloth patterns and body contours are blurred.

Firstly, our proposed T-WUTON architecture allows to improve the baseline CP-VTON. Indeed, our method generates spatial transformations of a much higher visual quality, which is specifically visible for stripes (1st row). It is able to preserve complex visual patterns of clothes and produces sharper images than CP-VTON and VTNFP on the edges. Secondly, we can observe the importance

of our distillation process with the visual results from S-WUTON. Since it has a non-masked image as input, it is able to preserve body details, especially the hands. Moreover, as shown in Fig. 1, S-WUTON is robust to a bad parsing and preserves a person’s attributes that are important for the virtual try-on task.

Generally, our method generates results of high visual quality while preserving the characteristics of the target cloth and of the person. However, VTNFP can surpass S-WUTON when models are crossing arms (4th row, right side), which is sometimes a failure case of our method. Note that this is not general, since on (3rd row, right side) and (4th row, left side) in Fig. 4 and on the two last columns in Fig. 5, models are crossing arms and S-WUTON manages to nicely compose the arms with the occluded cloth.

### 4.3 Quantitative results

Method	LPIPS	SSIM	IS	FID
Real data	0	1	3.135	0
CP-VTON on $p_{cp-vton}^*$	$0.182 \pm 0.049$	$0.679 \pm 0.073$	2.684	37.237
CP-VTON on $p_{t-wuton}^*$	$0.131 \pm 0.058$	$0.773 \pm 0.088$	2.938	16.843
T-WUTON	<b><math>0.101 \pm 0.047</math></b>	<b><math>0.799 \pm 0.089</math></b>	3.114	9.877
S-WUTON	NA	NA	<b>3.154</b>	<b>7.927</b>
VTNFP*	NA	0.803	2.784	NA
ACGPN*	NA	0.845	2.829	NA

**Table 1:** Quantitative results on paired setting (LPIPS and SSIM) and on unpaired setting (IS and FID). For LPIPS and FID, the lower is the better. For SSIM and IS, the higher is the better.  $\pm$  reports std. dev. The two last lines (methods with \*) are the results presented in ACGPN [41]. However, it has to be taken carefully since the experiments are performed on another dataset.

To further evaluate our method, we use four different metrics. Two are designed for the paired setting, that does not allow us to evaluate S-WUTON (because the input image is not masked), and one is for the unpaired setting. The first one for the paired setting is the linear perceptual image patch similarity (LPIPS) developed in [43], a state-of-the-art metric for comparing pairs of images. It is very similar to the perceptual loss we use in training (see Section 3.2) since the idea is to use the feature maps extracted by a pre-trained neural network to quantify the perceptual difference between two images. Different from the basic perceptual loss, they first unit-normalize each layer in the channel dimension and then learn a rescaling that match human perception.

Such as previous works, we also use the structural similarity (SSIM) [39] in the paired setting, inception score (IS) [33] and Fréchet Inception Distance (FID) [14] in the unpaired setting. We evaluate CP-VTON [38] on their agnostic person representation  $p_{cp-vton}^*$  (20 channels with RGB image of head and shape/pose information) and on  $p_{t-wuton}^*$ . Results are reported in Table 1.

#### 4.4 User study

We perform A/B tests on 7 users. Each one has to vote 100 times between CP-VTON and S-WUTON synthesized images, given reference person and target cloth. The user is asked to choose for the most realistic image, that preserves both person and target cloth details. The selected 100 images are a random subset of the test set in the unpaired setting. This subset is sampled for each user and is thus different for each user. There is no time limit for the users.

Let us denote  $p$  the probability that an image from S-WUTON is preferable to an image from CP-VTON. The users choose our method 88% of the time. In terms of statistical significance, it means that we can say  $p > 0.85$  with a confidence level of 98.7%.

ClothFlow and VTNFP also performed user studies where they compare to CP-VTON. The authors respectively report that users prefer their method 81.2% and 77.4% of the time. Note that the experiment was not performed in the same setting (dataset, number of users, number of pictures per user).

#### 4.5 Runtime analysis

	CP-VTON	VTNFP	ClothFlow	T-WUTON	S-WUTON
Parsing + pose	168ms	168ms	168ms	168ms	<b>0ms</b>
Try-on	<b>9ms</b>	>9ms	>0ms	13ms	13ms
Total	177ms	>177ms	>168ms	181ms	<b>13ms</b>

**Table 2:** Comparison of runtime of state-of-the-art architectures for virtual try-on. The time is computed on a NVIDIA Tesla V100 GPU.

In Table 2, we compare the runtime of our method to CP-VTON, ClothFlow and VTNFP. Note that the running times are estimated on a NVIDIA V100 GPU. For the human parsing and pose estimation networks, we use state-of-the-art models from [7,23]. These are based on shared neural backbones for the two tasks, which accelerates the computations.

The try-on architecture of T-WUTON and S-WUTON is slightly slower than that of CP-VTON, due to the non-shared encoder and the warping at each scale of the U-Net. However, with S-WUTON we remove the wall clock bottleneck of virtual try-on system, which is the human parsing and pose estimation. Doing so, we decrease by an order of magnitude the runtime of virtual try-on algorithms, from 6FPS to 77FPS.

We include comparisons with VTNFP and ClothFlow in the Table 2. Indeed, both models use human parsing and pose estimation. For VTNFP, they add a module on top of CP-VTON architecture, so their try-on architecture takes at least 9ms per image. For ClothFlow, the use of human parser and pose estimator gives a lower bound on the total runtime.

#### 4.6 The impact of adversarial loss in the teacher-student setting



**Fig. 5:** Visual comparison of the student model with and without the adversarial loss. Interestingly, the student model without the adversarial loss can not exploit information that is masked to the teacher, *e.g.* arms and hands.

We show the impact of the adversarial loss on S-WUTON. We train a variant student model S-WUTON without the adversarial loss. We provide a comparison of synthesized images in Fig. 5, and IS and FID scores in Table 3. The adversarial loss on the student model is a constraint to make the student model closer to the real data distribution and to not only follow the teacher’s distribution. Without the adversarial loss, the student model does not preserve person’s attributes, even though they are not masked.

	S-WUTON w/o adv.	S-WUTON
IS	2.912	<b>3.154</b>
FID	12.620	<b>7.927</b>

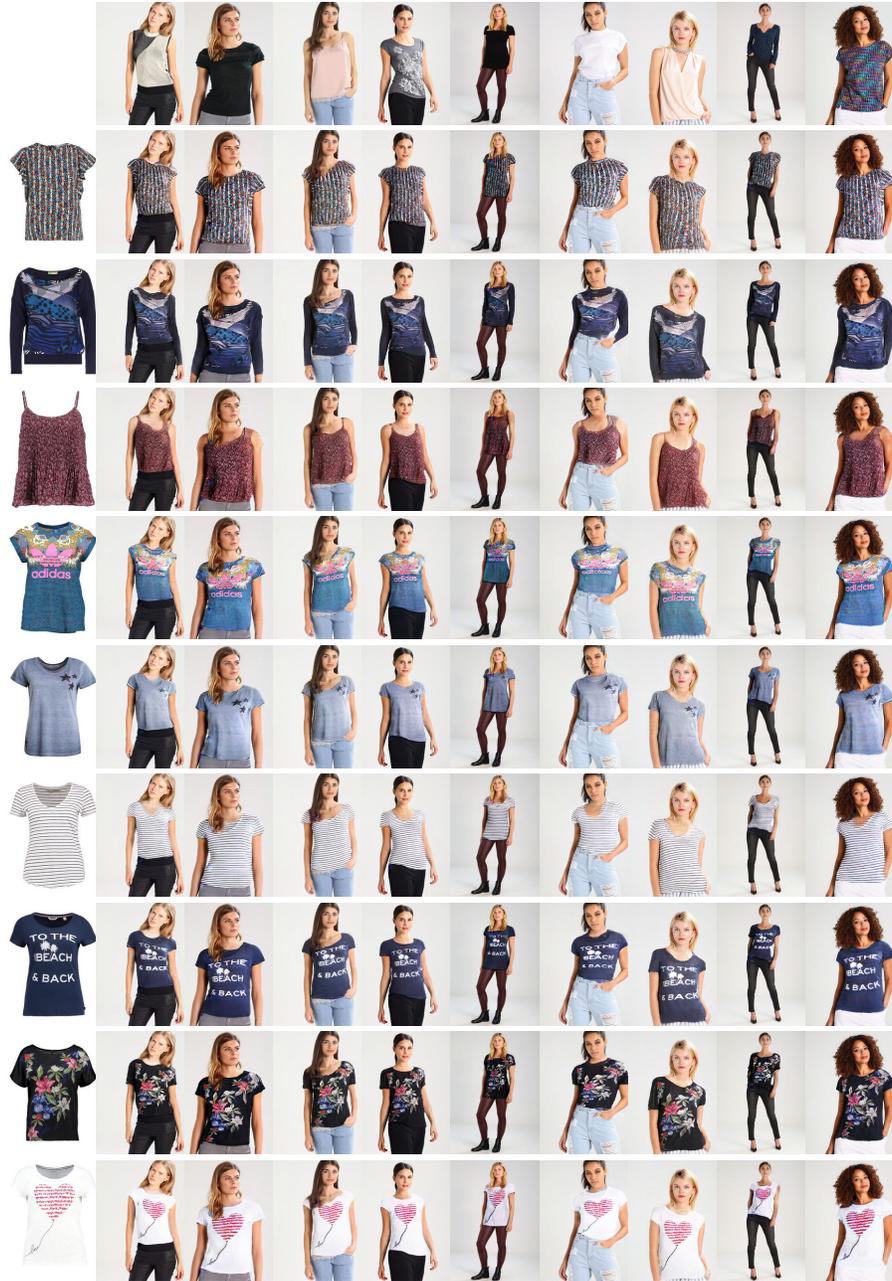
**Table 3:** Comparison of IS and FID scores of S-WUTON and S-WUTON without the adversarial loss.

## 5 Conclusion

In this work, we propose a teacher-student setting to distill the standard virtual try-on pipeline and refocus on the initial task: changing the cloth of a non-masked person. This leads to a significant computational speed-up and largely improves image quality. Importantly, this allows to preserve person’s attributes such as hands or accessories, which is necessary for a virtual try-on.

## Acknowledgements.

The authors thank David Picard and Marie-Morgane Paumard for their helpful advices.



**Fig. 6:** Our student model S-WUTON generates high-quality images and preserves both person's and cloth's attributes.

## References

1. Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018.
2. Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017.
3. Serge Belongie, Jitendra Malik, and Jan Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):509–522, 2002.
4. Fred L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on pattern analysis and machine intelligence*, 11(6):567–585, 1989.
5. Haoye Dong, Xiaodan Liang, Ke Gong, Hanjiang Lai, Jia Zhu, and Jian Yin. Soft-gated warping-gan for pose-guided person image synthesis. In *Advances in Neural Information Processing Systems*, pages 472–482, 2018.
6. Haoye Dong, Xiaodan Liang, Bochao Wang, Hanjiang Lai, Jia Zhu, and Jian Yin. Towards multi-pose guided virtual try-on network. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
7. Ke Gong, Xiaodan Liang, Dongyu Zhang, Xiaohui Shen, and Liang Lin. Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 932–940, 2017.
8. Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
9. Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. Drape: Dressing any person. *ACM Trans. Graph.*, 31(4):35–1, 2012.
10. Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017.
11. 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 (TOG)*, 33(4):105, 2014.
12. Xintong Han, Xiaojun Hu, Weilin Huang, and Matthew R. Scott. Clothflow: A flow-based model for clothed person generation. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
13. 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.
14. Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017.
15. Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
16. Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pages 448–456, 2015.

17. Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
18. Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015.
19. Nikolay Jetchev and Urs Bergmann. The conditional analogy gan: Swapping fashion articles on people images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2287–2292, 2017.
20. Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
21. Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard GAN. In *International Conference on Learning Representations*, 2019.
22. Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
23. Xiaodan Liang, Ke Gong, Xiaohui Shen, and Liang Lin. Look into person: Joint body parsing & pose estimation network and a new benchmark. *IEEE transactions on pattern analysis and machine intelligence*, 41(4):871–885, 2019.
24. Dominik Lorenz, Leonard Bereska, Timo Milbich, and Björn Ommer. Unsupervised part-based disentangling of object shape and appearance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
25. Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 99–108, 2018.
26. Youssef Alami Mejjati, Christian Richardt, James Tompkin, Darren Cosker, and Kwang In Kim. Unsupervised attention-guided image-to-image translation. In *Advances in Neural Information Processing Systems*, pages 3693–3703, 2018.
27. Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
28. Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J Black. Clothcap: Seamless 4d clothing capture and retargeting. *ACM Transactions on Graphics (TOG)*, 36(4):73, 2017.
29. Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *International Conference on Learning Representations*, 2016.
30. Amit Raj, Patsorn Sangkloy, Huiwen Chang, Jingwan Lu, Duygu Ceylan, and James Hays. Swapnet: Garment transfer in single view images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 666–682, 2018.
31. Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6148–6157, 2017.
32. 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.
33. Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
34. Rakshith R Shetty, Mario Fritz, and Bernt Schiele. Adversarial scene editing: Automatic object removal from weak supervision. In *Advances in Neural Information Processing Systems*, pages 7717–7727, 2018.

35. Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuilière, and Nicu Sebe. Deformable gans for pose-based human image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3408–3416, 2018.
36. Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
37. Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6924–6932, 2017.
38. Bochao Wang, Huabin Zheng, Xiaodan Liang, Yimin Chen, Liang Lin, and Meng Yang. Toward characteristic-preserving image-based virtual try-on network. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 589–604, 2018.
39. Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
40. Zhonghua Wu, Guosheng Lin, Qingyi Tao, and Jianfei Cai. M2e-try on net: Fashion from model to everyone. *arXiv preprint arXiv:1811.08599*, 2018.
41. 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.
42. Ruiyun Yu, Xiaoqi Wang, and Xiaohui Xie. Vtnfp: An image-based virtual try-on network with body and clothing feature preservation. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
43. Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018.
44. 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.
45. Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2223–2232, 2017.