



Cloth Interactive Transformer for Virtual Try-On

BIN REN, University of Trento and University of Pisa, Italy

HAO TANG, ETH Zürich, Switzerland

FANYANG MENG and DING RUNWEI, Peng Cheng Laboratory, China

PHILIP H. S. TORR, University of Oxford, UK

NICU SEBE, University of Trento, Italy

92

The 2D image-based virtual try-on has aroused increased interest from the multimedia and computer vision fields due to its enormous commercial value. Nevertheless, most existing image-based virtual try-on approaches directly combine the person-identity representation and the in-shop clothing items without taking their mutual correlations into consideration. Moreover, these methods are commonly established on pure convolutional neural networks (CNNs) architectures which are not simple to capture the long-range correlations among the input pixels. As a result, it generally results in inconsistent results. To alleviate these issues, in this article, we propose a novel two-stage cloth interactive transformer (CIT) method for the virtual try-on task. During the first stage, we design a CIT matching block, aiming at precisely capturing the long-range correlations between the cloth-agnostic person information and the in-shop cloth information. Consequently, it makes the warped in-shop clothing items look more natural in appearance. In the second stage, we put forth a CIT reasoning block for establishing global mutual interactive dependencies among person representation, the warped clothing item, and the corresponding warped cloth mask. The empirical results, based on mutual dependencies, demonstrate that the final try-on results are more realistic. Substantial empirical results on a public fashion dataset illustrate that the suggested CIT attains competitive virtual try-on performance.

CCS Concepts: • **Computing methodologies** → **Visual content-based indexing and retrieval**;

Additional Key Words and Phrases: Virtual try-on, transformer, garment transfer, cross attention

ACM Reference format:

Bin Ren, Hao Tang, Fanyang Meng, Ding Runwei, Philip H. S. Torr, and Nicu Sebe. 2023. Cloth Interactive Transformer for Virtual Try-On. *ACM Trans. Multimedia Comput. Commun. Appl.* 20, 4, Article 92 (December 2023), 20 pages.

<https://doi.org/10.1145/3617374>

1 INTRODUCTION

Virtual try-on (VTON), derived from fashion editing [7, 35, 60, 63], aims at transferring a desired in-shop clothing item onto a customer's body. If properly resolved, VTON will provide a time

This work was supported by the National Ph.D. in Artificial Intelligence for Society Program of Italy, the MUR PNRR project FAIR (PE00000013) funded by the NextGenerationEU and the EU H2020 AI4Media Project under Grant 951911.

Authors' addresses: B. Ren (Corresponding author), University of Trento & University of Pisa, via Sommarive 9, Trento, 38123, Italy; e-mail: bin.ren@unitn.it; H. Tang, ETH Zürich, Sternwartstrasse 7, Zürich, 8092, Switzerland; e-mail: hao.tang@vision.ee.ethz.ch; F. Meng and D. Runwei, Peng Cheng Laboratory, No. 2 Xingke 1st Street, Shenzhen Guangdong, 518000, China; e-mails: mengfy@pcl.ac.cn, dingrw@pcl.ac.cn; P. H. S. Torr, University of Oxford, Parks Road Oxford, OX1 3PJ, England; e-mail: philip.torr@eng.ox.ac.uk; N. Sebe, University of Trento, via Sommarive 9, Trento, 38123, Italy; e-mail: niculae.sebe@unitn.it.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2023 Copyright held by the owner/author(s).

1551-6857/2023/12-ART92

<https://doi.org/10.1145/3617374>

and energy-saving shopping experience in our everyday life. In practice, VTON has already been deployed in some big-brand clothing stores or e-commerce shopping applications owing to its convenience [17, 23, 30, 38].

However, most of the existing methods are designed based on 3D models pipelines [1, 20, 21, 29, 48] and follow the conventions of traditional computer graphics. Despite the detailed results, these methods require considerable labor resources, a significant amount of time investment, and complex data acquisition such as multi-view videos or 3D scans [40] that impede their widespread application. Alternatively, conditional **generative adversarial Networks (GANs)** based methods such as image-to-image translation or other image generation approaches [8, 12], recently made some positive progress. However, there remain some obvious artifacts in the generated results. To enhance the results of 2D image-based VTON methods, the classic two-stage pipeline VTON [22] was proposed, utilizing the first stage to warp the in-shop clothing item to a desired deformation style, and in the second stage, the warped cloth is aligned to the body shape of a given customer. While the visual result looks better than previous methods, there is still a significant gap between the overall visual quality and plausible generation. Many approaches following this pipeline, i.e., CP-VTON [55], ACGPN [59], and CP-VTON+ [37] were proposed with improved performance. However, these methods are limited to plain texture or simple-style clothes. Their performance suffers when dealing with complicated cases like rich textures or complex patterns. To address this issue, Xu et al. [58] introduced an intermediate operation that takes the transformation of the target person's image into consideration. But the improvement in visual performance is achieved at the cost of a more complex network architecture, which is time-consuming for training such a model. In addition, we also noticed that most of the previous methods rarely pay enough attention to the correlations between two crucial input information, i.e., the cloth-agnostic person information and the in-shop cloth information. Hence, there will be some inevitable mismatch phenomena occurring in the warped in-shop cloths, and consequently, it degrades the quality of the final try-on results. Moreover, for VTON, it's essential for a model to learn where to sample the pixels in the cloth image and where to reallocate them in the human body region. Hence, modeling the long-range dependencies is essential for achieving realistic try-on results. However, most of the previous methods use pure **convolutional neural networks (CNNs)** for modeling long-range dependencies. Since the pure CNN-based methods struggle to establish the long-range dependencies due to the design nature of the convolutional kernels, the final try-on performance is decreased.

Based on these observations, we assume that it would be advantageous to model the cloth-agnostic person information guided by the corresponding in-shop clothing information and vice versa. In addition, a better manner for capturing the long-range dependencies is also essential. To this end, based on the classic two-stage pipeline like VTON's [22], in this article, we propose a novel **Cloth Interactive Transformer (CIT)** method to address the aforementioned limitations. The overall architecture of the proposed CIT is depicted in Figure 1.

In the first stage (i.e., geometric matching stage), we design a CIT matching block that models long-range relations between the person and clothing representations interactively. Concurrently, a valuable correlation map is generated to boost the performance of the **thin-plate spline (TPS)** transform [4]. Unlike the traditional hand-crafted shape context matching strategies [34, 36, 46], which are only suitable for a certain feature type, the proposed CIT matching block (*Block-I*) has learnable features and can model long-range correlations via the cross-attention transformer encoders. As a result, the warped cloth becomes more natural and can fit a wearer's pose and shape more accurately. In the second stage (i.e., Try-On stage), unlike previous methods [37, 55] that treat the warped in-shop clothing item and its corresponding mask as a single input, we propose a novel CIT reasoning block (*Block-II*) that takes as input three distinct information, i.e., the cloth-agnostic person representations, the warped clothing item, and the mask of the warped clothing

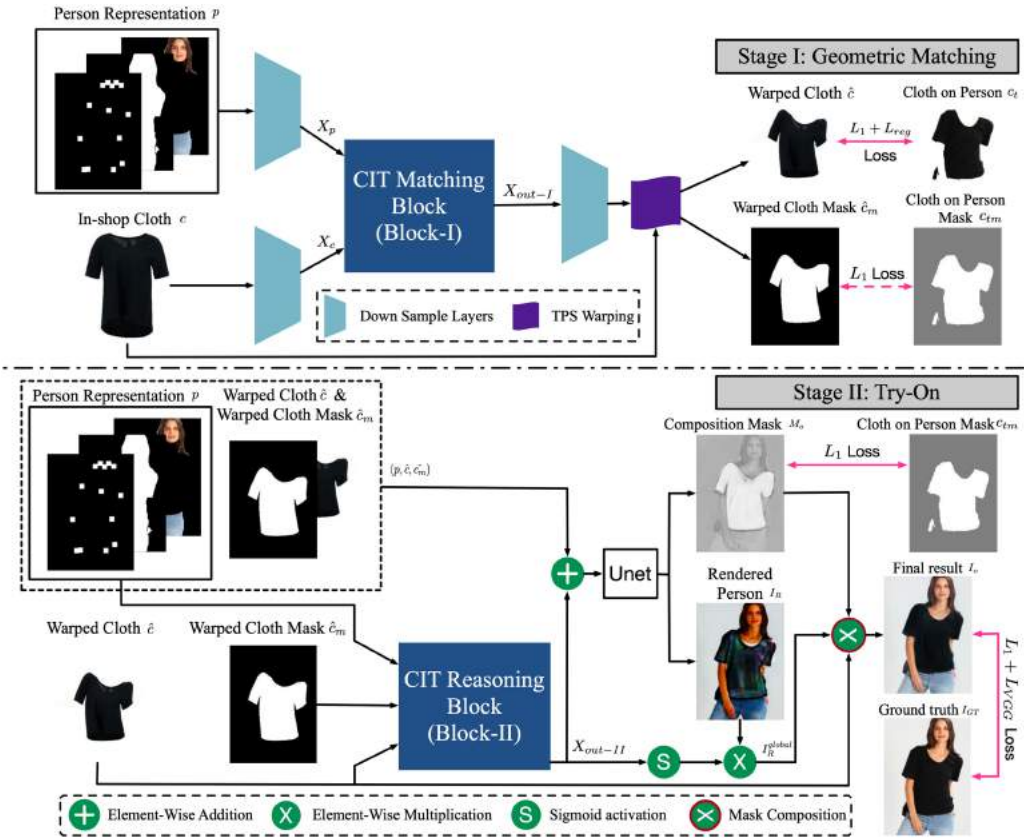


Fig. 1. The overall architecture of the proposed CIT for virtual try-on. The upper part is the Geometric Matching stage for warping the in-shop clothing items, while the bottom part is the Try-On stage for synthesizing the final try-on image of the person.

item. Through the CIT reasoning block, a more precise correlation among these three input data can be established, and this correlation is further utilized for strengthening the mask composition process. In addition, such a co-relation also serves as an attention map to activate the rendered person image, making the final results clearer and more realistic.

More specifically, in the CIT matching block, our primary objective is to improve the ability to model the person feature by encoding the target in-shop cloth feature. Since the in-shop clothing item is non-rigid, it is difficult to directly learn the matching relationship only from the clothing item. Hence, we resort to the correlation between the person and the target in-shop clothing item. With the help of this learned correlation, the person-related feature can be refined indirectly by the in-shop clothing features and vice versa. Such an analysis is also considered in the CIT reasoning block among the three input data. With the CIT matching and the CIT reasoning blocks, the correlation between both the person and in-shop cloth are updated synchronously and interact with each other.

In summary, our contributions are as follows:

- We design a novel two-stage (i.e., the geometric matching stage and the try-on stage) CIT method for the challenging virtual try-on task. CIT can well model the long-range

interactive relations between the cloth-agnostic person representations and the in-shop clothing items.

- We propose a new two-modality CIT matching block in the first geometric matching stage, making the in-shop clothing item can be better warped in the desired direction.
- We introduce a new three-modality CIT reasoning block in the second try-on stage. Based on this block, the more precise long-range correlations among three inputs (i.e., the cloth-agnostic person representation, the warped clothing item, and the corresponding mask of the warped cloth) can be well considered. As a result, CIT is able to obtain more realistic try-on image results.

2 RELATED WORK

Virtual Try-On (VTON), as one of the most popular tasks within the fashion area, has been widely studied by the research community due to its practical potential [2, 10, 15, 16, 18, 31]. Conventionally, this task was realized by computer graphic techniques, which build 3D models and render the output images via the precise control of geometric transformations or physical constraints [5, 6, 14, 20, 47, 61]. By using 3D measurements or representations, these methods can generate promising results for VTON, but the additional requirements such as 3D scanning equipment, computation resources, and heavy labor are not negligible.

Compared to 3D-based methods, 2D GANs based methods are more applicable to online shopping scenarios. Jetchev and Bergmann [26] proposed a conditional GAN to swap fashion articles with only 2D images. Another interesting GAN-based method SWAPGAN [32] solved the VITON task in an end-to-end manner, but it utilizes 3 generators, and the balance among all generators is hard to control. Gu et al. [19] proposed a GAN-based image transformation strategy that can automatically learn the mapping from a combination of pose and text to a target fashion image. However, this method didn't consider pose variations, it also required the availability of the paired images of both the in-shop clothes and the wearer during inference. Hence their applicability in practical scenarios is limited. Unlike the previous 3D-based or GAN-based methods, VITON [22] tackled this problem with a coarse-to-fine architecture, which first computed a shape context [3] **thin-plate spline (TPS)** transformation [4] for warping an in-shop clothing item on the target person and then blended the warped clothing item onto a given person. Note that the TPS transformation is a commonly used method of transforming a source image into a target image. It works by using a predefined set of constraints that define how the source image should be warped and transformed into the target image. The constraints are based on a set of surface points on the source image and their corresponding points on the target image. The warping algorithm then uses these points to calculate a set of transformations that will transform the source image into the target image. The transformations include scaling, rotation, translation, and shearing. The result is a warped version of the source image which accurately reflects the target image. This warping method used in the VTON task is to make the in-shop clothing item well warped toward the body shape of a given person. However, VITON [22] utilized hand-craft shape-context features for conducting the TPS transformation, which is not only time-consuming but also not robust when facing new samples. As an improvement, CP-VTON [55] and CP-VTON+ [37] adopted the learnable TPS transformation method proposed by [43] via a convolutional geometric matcher. Although the correlation between the person and the in-shape clothing features was established by such a differentiable TPS transformation, and the generated try-on results are better, there are still obvious artifacts when facing heavy occlusions, rich texture, or large transformations. ACGPN [59] was proposed to tackle these issues. Compared to CP-VTON, ACGPN used a semantic generation module to generate a semantic alignment of spatial layout. It also introduced the second-order difference constraint based on TPS. Though the performance is improved, the problem is still similar

to previous methods [22, 37, 55] because they didn't consider the global long-range interactive correlations between the person representation and the in-shop clothing item. Recently, Chopra et al. [9] proposed to solve the VTON task with a gated appearance flow. Though better results are achieved, the need of modeling the 3D geometric priors makes the overall procedure more complex.

To alleviate these problems, we propose a two-stage Cloth Interactive Transformer (CIT) method for the virtual try-on task. In particular, the proposed CIT can well capture the long-range dependencies in both stages. As a result, our method generates sharper and more realistic try-on images. **Long-Range Dependence Modeling.** Although CNN-based structures have shown excellent representation ability in various vision tasks such as classification, segmentation, and so on. The long-range dependencies are still hard to be established due to the limited receptive fields of the convolution kernels. For example, a convolutional kernel usually focuses on local neighbors (e.g., 3×3 or 5×5), while long-range relations would require the response at a position as a weighted sum of the features at all other positions. This limitation raises huge challenges for many applications where long-range relationships are needed.

To overcome this limitation, the attention mechanism [50, 51, 54] has been widely used in vision tasks with CNN architectures though it was initially designed for natural language processing tasks. In addition, non-local neural networks [56] were designed based on the self-attention mechanism, allowing the model to capture the long-distance dependencies in the feature maps. However, this approach suffers from high memory and computation costs. [45] proposed an attention gate model to increase the sensitivity of a base model. Besides, the multilayer perceptrons (MLP) are also proposed for modeling the long-range relation, but it may heavily affect the efficiency [42, 52]. Moreover, Transformer [54] was first introduced for neural machine translation tasks because it can model long-range dependencies in sequence-to-sequence tasks and captures the relations between arbitrary positions in the given sequence. Unlike previous CNN-based methods, Transformers are built solely on self-attention operations, which are strong in modeling the global context. After Transformers demonstrated their overwhelming power on a broad range of language tasks (e.g., text classification, machine translation, or question answering [11, 28, 39, 41, 49]). Recently, Transformer-based frameworks have also shown their effectiveness on various vision tasks [33]. In particular, **vision Transformer (ViT)** [13, 53] splits the image into patches and models the correlation between these patches as sequences, and then the core self-attention module of ViT is stacked for modeling the long-range dependencies.

To this end, in this article, we also utilize a vision Transformer for handling the long-range dependencies among the cloth-agnostic person representation, the in-shop clothing item (for both the original cloth and the warped cloth), and the corresponding mask of the warped clothing item in a novel cross-modal manner.

3 CLOTH INTERACTIVE TRANSFORMER

In this section, we first give an overall introduction and the necessary notations of the proposed CIT method for virtual try-on in 3.1. Then we introduce the core modules (i.e., the Interactive Transformer I and II) of CIT in Section 3.2. Based on the interactive transformer -I and -II, we show the details of the CIT matching block (Block-I) and the CIT reasoning block (Block-II) in Section 3.3 and Section 3.4, respectively. Finally, the optimization objectives of the proposed CIT for both stages are described in detail in Section 3.5.

3.1 Overview and Notations

For the 2D image-based VTON task, the target in-shop clothing item is different from the source clothing item that is worn by a given person. Specifically, given a person image $I \in \mathbb{R}^{3 \times h \times w}$ and one

in-shop clothing image $c \in \mathbb{R}^{3 \times h \times w}$. Our goal is to generate the image $I_o \in \mathbb{R}^{3 \times h \times w}$ where a person I wears the cloth c . Hence, what we need to do first is to reduce the side effects of the source clothes, like color, texture, or shape. Meanwhile, it's also necessary to preserve the information about the given person as much as possible, including the person's face, hair, body shape, and pose. To this end, we adopt the same pipeline as [22] for the person representation p from I . It contains three components, the 18-channel feature maps for the human pose, the 1-channel feature map for the body shape, and the 3-channel RGB image. Note that the RGB image contains only the reserved regions of a person (i.e., face, hair, and lower parts of the person) for maintaining the identity of this person.

The basic structure of the proposed CIT is in a two-stage (i.e., see the geometric matching stage and the try-on stage in Figure 1) pipeline, which is also adopted by CP-VTON [55] and CP-VTON+ [37]. In particular, the former takes as input the cloth-agnostic person representation p and an in-shop clothing item c to produce a warped cloth \hat{c} and a warped mask \hat{c}_m based on the given person's pose and shape. The latter uses the warped cloth \hat{c} , the corresponding warped mask \hat{c}_m together with the person representation p to generate the final person image with the worn in-shop cloth. In the first geometric matching stage, we propose a CIT matching block (Block-I, see the upper part in Figure 2 for the details), which takes the person feature X_p and the in-shop cloth feature X_c as inputs. Then X_p and X_c are generated by two similar feature extractors from p and c , respectively (see the first geometric matching stage in Figure 1). After that, it produces a correlation feature X_{out-I} followed by a down-sample layer for regressing the parameter θ . Note that θ is used for warping the original in-shop clothing c to the target on-body style \hat{c} via an interpolation method named the **thin-plate spline (TPS)** warping module [43]. Specifically, given two images with some corresponding control points in different positions, these control points can be well aligned from one (i.e., in the in-shop clothing item) to another (i.e., in the corresponding human body region) with the thin plate spline interpolation operation in a geometry estimation manner (i.e., local descriptor extraction, descriptor matching, transformation-related parameter estimation) [22].

In addition, the TPS operation we adopted in this article is the same as the one used in CP-VTON [55] from [43]. It first utilizes its differentiable modules to conduct a transformation by mimicking the geometry estimation procedure in a learnable manner from c to the \hat{c} . Meanwhile, the corresponding mask \hat{c}_m of \hat{c} is also produced based on θ via TPS warping operation. In the second stage, we utilize the warped cloth \hat{c} and the warped mask \hat{c}_m together with the person representation p as inputs of the CIT reasoning block (Block-II, see the bottom part in Figure 2 for the details). And the output X_{out-II} of the CIT reasoning block is used to guide the final mask composition for generating more realistic try-on results.

3.2 Interactive Transformer

Having leveraged the self-attention mechanism, Transformer is capable of modeling long-range dependencies. Given this inherent ability of the Transformer, we propose the innovative Interactive-Transformer for exploring the correlation between the person and the clothing item in the VTON task. There are two types of Interactive Transformers in the proposed CIT. The first version, i.e., Interactive Transformer I, is employed in the first geometric matching stage. The second version, i.e., Interactive Transformer II, is utilized in the second try-on stage. They are based on the basic Transformer encoders and the cross-modal Transformer encoders, and their detailed description is depicted in Figure 2.

Regarding a standard Transformer encoder, a positional embedding is initially added to the input feature as elucidated in [54]. This is beneficial for keeping the initial spatial relation of the input. After the position embedding, the input feature will be projected into queries Q_m , keys K_m , and

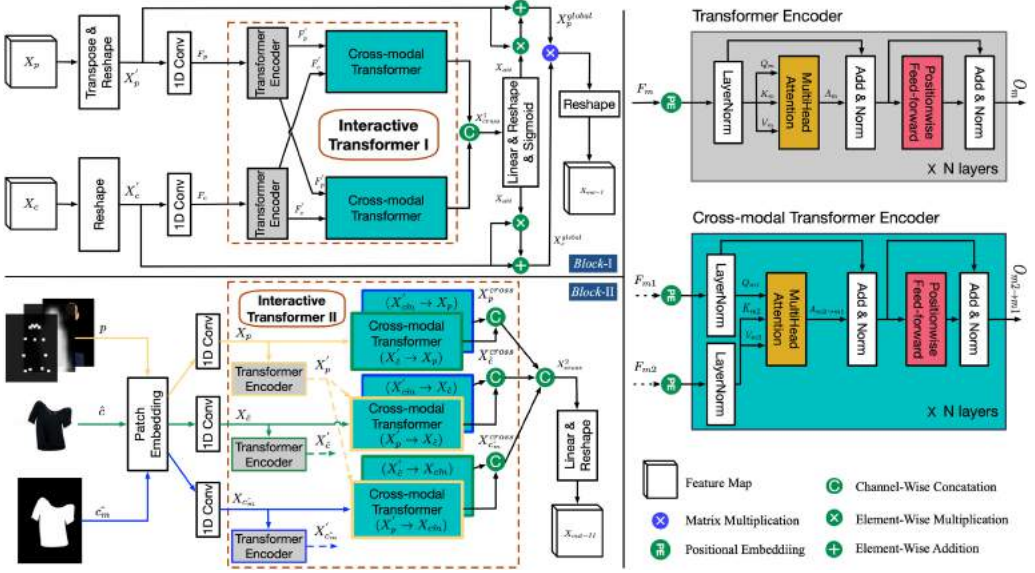


Fig. 2. The key components of the proposed CIT for virtual try-on. The upper area on the left is the CIT Matching Block (*Block-I*), while the bottom area on the left indicates the CIT Reasoning Block (*Block-II*). On the right, the normal Transformer encoder and the proposed cross-modal Transformer encoder are shown in detail.

values V_m by a linear layer. Subsequently, the output of the attention layer A_m is derived as

$$A_m = \text{softmax} \left(\frac{Q_m K_m^T}{\sqrt{d}} \right) V_m, \quad (1)$$

where d is the dimension of the query, key.

The aforementioned self-attention mechanism is usually employed for only one type of input data. However, in the two-stage VTON task, to accurately capture a more precise match between the information of the person and the cloth information, there are several pairs of correlations we can't overlook. Notably, in the geometric matching stage, we need to consider the correlation between the cloth-agnostic person representation p and the in-shop clothing item c since such a correlation is indispensable for producing a reasonable warped cloth \hat{c} . In the second try-on stage, there are three types of inputs i.e., p , \hat{c} , as well as \hat{c}_m . To proficiently model the long-range connection of each two of them (i.e., p and \hat{c} , p and \hat{c}_m , as well as \hat{c} and \hat{c}_m) is also a crucial issue since a well-captured correlation usually yields a good match between the person's body and the in-shop cloth.

Based on such an observation, instead of using only the self-attention layer in a Transformer encoder for processing a single-modal input, we propose a cross-modal Transformer encoder based on a cross-attention mechanism. Note that we treat each kind of input as a certain single-modal input since each of them provides a specific input. For example, p is for person identity, c and \hat{c} correspond to the texture, and \hat{c}_m is related to the shape information. And the cross-attention is computed as follows:

$$A_{m2 \rightarrow m1} = \text{softmax} \left(\frac{Q_{m1} K_{m2}^T}{\sqrt{d}} \right) V_{m2}, \quad (2)$$

where we adopt the first input (i.e., person representation p) as query Q_{m1} , and the second input (i.e., the in-shop clothing item c) as the keys K_{m2} and values V_{m2} . Based on such a cross-interactive manner, each kind of input keeps updating its sequence via the external information from the multi-head cross-attention module. As a result, one modality will be transformed into a different set of key/value pairs to interact with another modality.

Interactive Transformer I is shown in the red dash box of the upper area of Figure 2. It consists of two regular Transformer encoders (depicted in gray) and two cross-modal Transformer encoders (depicted in light blue) that are directly applied to feature maps. We use $selfTrans(\cdot)$ and $crossTrans(\cdot)$ to indicate the operators of these two kinds of Transformer encoders. Given two input features F_p and F_c with dimension (C, B, S) . Note that the dimension of F_p and F_c are reshaped from input features X_p and X_c with original dimension (B, C, H, W) . Here B, C, H, W denote the batch size, the number of channels, the height, and the width of the input features X_p and X_c , $S = H \times W$ denotes the spatial dimension. Then each of them will go through its corresponding N -layer regular Transformer encoder first. After that we get the processed features F'_p and F'_c as follows:

$$\begin{aligned} F'_p &= selfTrans(F_p), \\ F'_c &= selfTrans(F_c). \end{aligned} \quad (3)$$

Then the cross-modal Transformer encoder is used for modeling the cross-modal long-range correlations between F'_p and F'_c :

$$X_{cross}^1 = cat(crossTrans(F'_p, F'_c), crossTrans(F'_c, F'_p)), \quad (4)$$

here $crossTrans(X'_p, X'_c)$ indicates that we utilize X'_c as the keys and values while we use X'_p as the queries. On the other hand, $crossTrans(X'_c, X'_p)$ indicates that the keys and values are coming from X'_p and queries come from X'_c . After concatenating the outputs from the two cross-modal Transformer encoders, we get the output X_{cross}^1 of the Interactive Transformer I. It can strengthen the correlation matching ability.

Interactive Transformer II is shown in the red dash box at the bottom area of Figure 2. Similar to the Interactive-Transformer I, the Interactive Transformer II is also constructed by combining regular Transformer encoders and cross-modal Transformer encoders.

The Interactive Transformer II is designed mainly for exploring the correlations between every two inputs among three total inputs (i.e., p, \hat{c} , and \hat{c}_m). In particular, we adopt 3 regular Transformer encoders and 6 cross-modal Transformer encoders for constructing the Interactive Transformer II. Note that for better illustration, we depict $X_p, X_{\hat{c}}$, as well as $X_{\hat{c}_m}$ and their corresponding information flows in yellow, green, and blue, respectively.

Within the Interactive Transformer II, there are three input features i.e., $X_p, X_{\hat{c}}$, and $X_{\hat{c}_m}$. Each of them works as the Query element within its own branch while working as the Key and Value elements in the other two branches. Specifically, we take the feature X_p (depicted in yellow in Figure 2) that comes from person representation as the detailed introduction. Once we get X_p after the 1D convolutional layer that is out of the red dash box. There are two pathways for X_p to pass through. The first one is to directly let it meet two cross-modal Transformer encoders (i.e., the green-border cross-modal Transformer encoder between $X'_{\hat{c}}$ and X_p , as well as the blue-border cross-modal Transformer encoder between $X'_{\hat{c}_m}$ and X_p). Another one is to let X_p pass through a regular Transformer encoder for producing the updated feature X'_p . Note that here $(X'_{\hat{c}} \rightarrow X_p)$ within the green-border cross-modal Transformer encoder means we utilize X_p as Query and $X'_{\hat{c}}$ as Key and Value, while $(X'_{\hat{c}_m} \rightarrow X_p)$ within the blue-border cross-modal Transformer encoder

indicates we use X_p as Query and X'_{c_m} as Key and Value. X'_c and X'_{c_m} are the updated features from X_c and X_{c_m} after their corresponding regular Transformer encoders. We formulate such procedures of the first yellow branch as follows:

$$X_p^{cross} = \text{cat}(\text{crossTrans}(X'_p, X'_c), \text{crossTrans}(X'_p, X'_{c_m})). \quad (5)$$

Similarly, we also get the output of the middle green branch X_c^{cross} and the output of the bottom blue branch $X_{c_m}^{cross}$. Finally, the overall output of the Interactive Transformer II is

$$X_{cross}^2 = \text{cat}(X_p^{cross}, X_c^{cross}, X_{c_m}^{cross}). \quad (6)$$

3.3 CIT Matching Block

Based on our Interactive Transformer I, we propose the CIT Matching block (Block-I) to boost the performance of the TPS transformation by strengthening the long-range correlation between $X_p \in \mathbb{R}^{B \times C \times H \times W}$ and $X_c \in \mathbb{R}^{B \times C \times H \times W}$. Here B , C , H , and W indicate batch size, channel number, and the height and width of a given feature. To utilize the Transformer encoder for modeling long-range dependencies, we first adjust the dimensions of X_p and X_c from (B, C, H, W) to (B, C, S) , forming the X'_p and X'_c . Note that $S = H \times W$. Besides, a 1D convolutional layer is also adopted to ensure that each element of each input sequence can get sufficient awareness of its neighborhood elements. When we get F_p and F_c after the convolutional layers, the proposed Interactive Transformer I is applied to F_p and F_c for capturing the long-range correlation between the person-related and the in-shop cloth-related features. As a result, we get the result (i.e., X_{cross}^1) of the proposed CIT matching block. These procedures are depicted in Figure 2 with detailed annotations.

Instead of directly adding this long-range relation to features X_p or X_c , we strengthen each of them by a global strengthened attention X_{att} operation as follows:

$$X_{(\cdot)}^{global} = X_{(\cdot)} + X_{(\cdot)} \times X_{att}, \quad (7)$$

Here \times means an element-wise multiplication, (\cdot) indicates that both features X_p and X_c follow the same form. Note that X_{att} is produced from X_{cross}^1 by a linear projection and a sigmoid activation. Based on this operation, the element position relation of each input will be activated by the sigmoid activation function. In particular, when it is applied to the input feature as attention, both the position information of each element within each input and the correlation between two inputs can be kept in a balanced manner. Then a matrix multiplication between X_p^{global} and X_c^{global} is conducted. The output X_{out-I} of the proposed CIT matching block is finally obtained after a reshape operation, which represents the improved correlation between the person and clothing features. These procedures can be defined as follows:

$$x_{out-I} = \text{Reshape}((X_c^{global})^T \times X_p^{global}). \quad (8)$$

Here X_p^{global} and X_c^{global} have the same dimension (B, C, S) , and the output X_{out-I} is in dimension (B, S, H, W) .

3.4 CIT Reasoning Block

Previous works (i.e., CP-VTON [55] and CP-VTON+ [37]) first concatenate the person information p , the warped cloth information \hat{c} , and the warped clothing mask \hat{c}_m together. Then the concatenated input is directly sent to one UNet model as a single input to generate a composition mask M_o as well as a rendered person image I_R . However, such a rough concatenate operation may lead to coarse information matching, and consequently, it would be difficult to achieve a well-matched final try-on result.

To this end, we propose the CIT Reasoning block (Block-II) depicted in Figure 2, aiming at modeling such more complicated correlations among p , \hat{c} , and \hat{c}_m . Firstly, we adopt the patch embedding operation [13] to all these three inputs. Then each of them goes through a 1D convolutional layer to ensure the relation modeling of each element with its neighbor elements. After that, we get X_p , $X_{\hat{c}}$, and $X_{\hat{c}_m}$. To well capture the complicated long-range correlations among these features, we apply the proposed Interactive Transformer II to X_p , $X_{\hat{c}}$, as well as $X_{\hat{c}_m}$. Then the output X_{out-II} of Interactive Transformer II is utilized to guide the final mask composition for a better generation as follows:

$$\begin{aligned} I_R^{global} &= \text{sigmoid}(X_{out-II}) \times I_R, \\ I_o &= M_o \times \hat{c} + (1 - M_o) \times I_R^{global}, \end{aligned} \quad (9)$$

here *sigmoid* indicates the Sigmod activation function.

3.5 Optimization Objectives

The first stage of CIT is trained with sample triplets (p, c, c_m) , while the second stage is trained with (p, \hat{c}, \hat{c}_m) . In addition, in the first matching stage, we adopt the same optimization objectives as CP-VTON+ [37]:

$$\mathcal{L}_{Matching} = \mathcal{L}_1(\hat{c}, c_t) + \frac{1}{2}\mathcal{L}_{reg}, \quad (10)$$

where \mathcal{L}_1 indicates the pixel-wise L1 loss between the warped result \hat{c} and the ground truth c_t . \mathcal{L}_{reg} indicates the grid regularization loss, and it can be formalized as follows:

$$\mathcal{L}_{reg}(G_x, G_y) = \sum_{i=-1,1} \sum_x \sum_y |G_x(x+i, y) - G_x(x, y)| + \sum_{j=-1,1} \sum_x \sum_y |G_y(x, y+j) - G_y(x, y)|, \quad (11)$$

where G_x and G_y indicate the grid coordinates of the generated images along the x and the y directions.

In the second stage, the optimization objective is as follows:

$$\mathcal{L}_{Try-on} = \|I_o - I_{GT}\|_1 + LVGG + \|M_o - c_{tm}\|_1. \quad (12)$$

The first item aims at minimizing the discrepancy between the output I_o and the ground truth I_{GT} . The second item, the VGG perceptual loss [27], is a widely used loss item in image generation tasks. It is an alternative to pixel-wise losses and it attempts to be closer to the perceptual similarity of human beings. The VGG loss is based on the ReLU activation layers of the pre-trained 19-layer VGG network. It can be expressed as follows:

$$\mathcal{L}_{VGG} = \frac{1}{W_{i,j}H_{i,j}} \sum_{x=1}^{W_{i,j}} \sum_{y=1}^{H_{i,j}} (\phi_{i,j}(I_{GT})_{x,y} - \phi_{i,j}(I_o)_{x,y})^2, \quad (13)$$

where $W_{i,j}$ and $H_{i,j}$ describe the dimensions of the respective feature maps within the VGG network. $\phi_{i,j}$ indicates the feature map obtained by the j th convolution before the i th max pooling layer within the VGG19 network. The third item is used to encourage the composition mask M_o to select the most suitable warped clothing mask c_{tm} as much as possible.

4 EXPERIMENTS

Datasets. We conduct all the experiments on the same dataset collected by Han et al. [22] that used in VITON [22], CP-VTON [55], and CP-VTON+ [37]. Note that due to copyright issues, we only use the reorganized version as the previous works [37, 55] did. It contains around 19,000 front-view women and top clothing image pairs. Specifically, there are 16,253 cleaned pairs which are split into a training set and a validation set with 14,221 and 2,032 pairs, respectively. In the

training set, the target cloth and the cloth worn by the wearer are the same. However, in the test stage, *there are two kinds of test settings*. The first one is the same as the training settings, where the target clothing item and the clothing item worn by the wearer are the same (we refer to this case as a **retry-on setting** because it is just like the wearer takes off the cloth first then retrieves this cloth on, hence we have the ground truth for this case). Another setting means that the target clothing item is different from the one worn by the wearer (we refer to this kind as the **try-on setting**).

Evaluation Metrics. To evaluate the performance of our method. We firstly adopt the Jaccard Score [25] for the **retry-on** case (i.e., with ground truth) in the first stage. We also follow [24, 37, 59] to use the **Structural Similarity (SSIM)** [57], **Learned Perceptual Image Patch Similarity (LPIPS)** [62], **Peak Signal-to-Noise Ratio (PSNR)**, **Frechet Inception Distance (FID)**, and **Kernel Inception Distance (KID)** metrics in the second stage. Note that we adopt the original human image with the original clothing item as the reference image for SSIM and LPIPS (the lower, the better), and the parsed segmentation area for the current upper clothing is used as the reference for calculating the JS score. For **try-on** case (no ground truth), we evaluate the performance of our method and other state-of-the-art methods by the **Inception Score (IS)** [44].

Implementation Details. For the geometric matching stage, we build 2 feature extractors (see the downsample layers shown in Figure 1) with 4 2-strided downsampling convolutional layers followed by 2 1-strided ones (filter numbers: 64, 128, 256, 512, 512) to generate X_p and X_c . Those two extractors are only different in input channels. The regression networks before the TPS warping operation contain 2 2-strided convolutional layers, 2 1-strided ones (filter numbers: 512, 256, 128, 64), and 1 FC layer with output size 50. For the try-on stage, the U-net used here consists of 6 2-strided down-sampling convolutional layers (filter numbers: 64, 128, 256, 512, 512, 512) and 6 up-sampling (filter numbers: 512, 512, 256, 128, 64, 4). Each convolutional layer is followed by one InstanceNorm layer, and a Leaky Relu with slope is set to 0.2. Note that we stack 3 CIT encoders in both matching and reasoning blocks.

Our training settings are similar to CP-VTON and CP-VTON+. Both stages are trained for 200K steps with batch size 4. Moreover, for Adam optimizer, β_1 and β_2 are set to 0.5 and 0.999, respectively. The learning rate was firstly fixed at 0.0001 for the first 100K steps and then linearly decayed to zero for the rest of the steps. All input images are resized to 256×192 , and the output images have the same resolution. The source code and trained models are available at <https://github.com/Amazingren/CIT>.

4.1 Qualitative Comparisons

To validate the performance of the proposed CIT for virtual try-on, we first present the visualization results of both stages, including the warped clothing items and the final try-on person images.

Comparison of Warping Results. We visualize the warped results of clothes for both retry-on and try-on settings in Figure 3. Note that in this article, our visualization method is similar to CP-VTON [55], and CP-VTON+ [37]. For ACGPN [59], we directly adopt its officially released code, and it produced the gray-background results. Figure 3 shows that the proposed CIT generates sharper and more realistic warped clothing items than the other methods like CP-VTON, CP-VTON+, and ACGPN. This typically happens for texture-rich cases, such as the case with line stripes (see the last row of the same-pair case and the second row of the different-pair case), or in the presence of logos (see the 2nd row of the same-pair case and the third row of the different-pair case), and so on. We marked out these obvious artifacts from other methods with red dash boxes in Figure 3.

Comparison of Try-On Results. Figure 4 shows the try-on results. We can see that the proposed CIT outperforms other methods. Specifically, the proposed CIT can keep the original clothing



Fig. 3. Qualitative comparisons of the warped cloths by the proposed CIT-based geometric matching stage. The left is for the retry-on setting (i.e., in the same cloth) while the right is for the try-on setting (i.e., in different clothing items).

Table 1. User Study Comparison on Two Questions

Method	Q1	Q2
CP-VTON [55]	19.5	14.6
CP-VTON+ [37]	24.8	25.2
ACGPN [59]	23.6	24.8
CIT (Ours)	32.1	35.4

Q1 denotes 'Which image is the most photo-realistic?', and Q2 denotes 'Which image preserves the details of the in-shop clothing item the most?' in the user study.

texture and its pattern as much as possible, and the final results are more realistic and natural. Compared to our method, other approaches display many artifacts, for example, the irregular logo pattern (the 1st row), the over-warped cloth texture (the 3rd row), and the ridiculous results for unique or complicated-style cloth (the last row). We also marked out these artifacts in red dash boxes in Figure 4. In addition, more qualitative try-on results of the proposed method can be found in Figure 5.

User Study. We also evaluate the proposed CIT and other methods via a user study. We randomly select 120 sets of reference and target clothing images from the test dataset. Given the reference images and the target in-shop clothing items, 30 users are asked to choose the best outputs of our model and baselines (i.e., CP-VTON, CP-VTON+, and ACGPN) according to the two questions: (Q1) Which image is the most photo-realistic? (Q2) Which image preserves better the details of the target clothing? As shown in Table 1, we can see that the proposed CIT achieves significantly better results than the other methods, which further demonstrates that our model generates more realistic images. And CIT also preserves the details of the clothing items compared to the other methods.



Fig. 4. Qualitative comparisons of different state-of-the-art methods.

4.2 Quantitative Evaluation

To further evaluate the performance of our CIT, we adopt five evaluation metrics, i.e., JS, SSIM, LPIPS, PSNR, IS, FID, and KID for numerical comparison. JS is to evaluate the quality of the warped mask in the first geometric matching stage with same-pair test samples, which is equivalent to the IoU metric used in CP-VTON+ but is more convenient for implementation. Note that we take cloth masks of the person as a reference image. Other metrics are designed to evaluate the performance of the second try-on stage.

The results of JS are shown in Table 2. Though our CIT doesn't achieve the best JS score, our visual results presented in Figure 3 are the most reasonable ones. We think with the help of the proposed Interactive Transformer I in the CIT matching block, our method can learn more reasonable texture transformation patterns. This learned strong texture-focused transformation pattern might affect the shape alignment. This may come from the reason that the JS score only focuses on the shape alignment aspect between the ground truth mask and the warped clothing mask. Consequently, it can not well reveal the overall quality of the final generated human images. For



Fig. 5. More qualitative results of the proposed method.

instance, though Table 2 shows that CP-VTON+ has the best JS score of 0.812, which is higher than ours (0.800), the qualitative results show that our method is superior to CP-VTON+. So the shape-only related evaluation metrics i.e., JS or IoU, are not always indicating a better overall visual result. We also conducted ablation experiments in the following to support this conclusion (see the comparison results in Table 3 between B3 and B4). In addition, in terms of the realism of the generated images, the proposed CIT can also achieve the best KID and the second-best results, which confirm the effectiveness of our method.

Table 2. Quantitative Comparison in Terms of JS, SSIM, LPIPS, PSNR, IS, FID, and KID Evaluation Metrics

Method	JS \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	IS \uparrow	FID \downarrow	KID \downarrow
CP-VTON [55]	0.759	0.800	0.126	14.544	2.832	35.16	2.245
CP-VTON+ [37]	0.812	0.817	<u>0.117</u>	21.789	3.074	25.19	1.586
ACGPN [59]	-	0.846	0.121	<u>23.080</u>	2.924	13.79	<u>0.818</u>
CIT (Ours)	<u>0.800</u>	<u>0.827</u>	0.115	23.464	3.060	<u>13.97</u>	0.761

For JS, SSIM, PSNR, and IS, the higher, the better, while for LPIPS, FID, and KID, the lower, the better. What's more, IS, FID, and KID are used to evaluate the unpaired **try-on** setting while the rest are all for paired **retry-on** setting. Note that the optimum results are **bolded** while the second-best results are underlined.

Table 3. Ablation Studies of the Proposed CIT for Virtual Try-on

Method	JS \uparrow	SSIM \uparrow	LPIPS \downarrow	IS \uparrow	FID \downarrow	KID \downarrow	Q1	Q2
B0 [37]	0.812	0.817	0.117	3.074	25.19	1.586	-	-
B1 (CIT Matching only)	0.800	0.808	0.123	3.020	14.76	0.779	-	-
B2 (CIT Reasoning only)	0.812	0.821	0.125	3.105	14.87	0.784	-	-
B3 (Full: B1+B2)	0.800	0.827	0.115	3.060	13.97	0.761	70.8	69.5
B4 (Full + L_1 mask loss)	0.813	0.829	0.110	3.005	14.32	0.788	29.2	30.5

For the **retry-on** setting, we adopt SSIM, PSNR, and LPIPS to evaluate the performance. The numerical results are shown in Table 2. It can be seen that the proposed CIT achieves the best numerical evaluation results on SSIM and LPIPS compared to others. For the **try-on** setting, we use IS for evaluation. The results in Table 2 show that our CIT achieves just a slightly lower IS score of 3.060 compared to 3.074 of CP-VTON+. We think the most possible reason for this phenomenon is that IS is an objective metric that usually was used to measure the quality of the generated images at the feature level based on image diversity and clarity. Hence, it may ignore some pixel-level properties.

Overall, though we do not obtain the best quantitative scores on JS and IS metrics, our proposed CIT can generate sharper and more realistic try-on images compared to others. For ACGPN, we also test the performance based on the corresponding official checkpoints. However, because the test set of ACGPN is different from others, so we only present its visual results in Figure 3 and Figure 4.

4.3 Ablation Study and Discussion

To validate the effectiveness of each part of the proposed CIT, we conduct four ablation experiments (i.e., B1, B2, B3, and B4 in Table 3). CP-VTON+ [37] is adopted as the baseline (B0) of this article. B1 means that we only use the proposed CIT matching block in the first geometric matching stage but keep the second stage the same as B0; B2 means that we only use the proposed CIT reasoning block in the second try-on stage while keeping the first stage the same as B0; B3 is the final version adopted in this article, which contains both the proposed CIT matching and reasoning blocks; B4 is built based on B3, but with an extra L_1 loss item for providing more stricter constraints between the generated warped clothing mask \hat{c}_m and the ground truth cloth mask c_{tm} of the given person (depicted with red dash lines in Figure 1). The setting of this experiment is designed to support the conclusion that a higher JS (or IoU) score doesn't indicate more enjoyable visual results since it ignores the texture or pattern aspect of the overall quality. The overall matching loss of B4 can be summarized as follows:

$$L_{Matching} = L_1(\hat{c}, c_t) + L_1(\hat{c}_m, c_{tm}) + \frac{1}{2} \cdot L_{reg}, \quad (14)$$



Fig. 6. Qualitative comparisons of ablation studies between B3 and B4.

From the comparison between B0 [37] and B1 in Table 3, though CP-VTON+ achieves a better JS score, the qualitative results presented in Figure 3 indicate that B1 can generate more reasonable and natural warped clothes (for both the retry-on and try-on settings). In other words, the proposed CIT Matching block can well capture more texture-related latent patterns with the help of the proposed Interactive Transformer I. The comparison between B0 and B2 shows that the numerical results, i.e., SSIM, and IS, are improved when we apply the proposed CIT reasoning block to the warped results from B0. This demonstrates that the proposed CIT reasoning block is more effective in generating the try-on results. B3, the combination of B1 and B2, is the final version of the proposed CIT. It produces not only the more natural warped clothing items but also achieves the more realistic try-on result. Hence, we think that for the two-stage 2D image-based VTON task, the JS or IoU metric focuses on only one aspect (i.e., shape) of the overall quality. Hence, the final try-on results are not always better when JS or IoU scores are higher. To support the above conclusion, we design B4 as a supplementary experiment based on B3. In Table 3, B4 obtains nearly all the best numerical results except the IS score. However, the visual comparison in Figure 6 shows that its virtual try-on results are far from satisfactory compared to B3. In addition, we asked 15 users for



Fig. 7. Several failure cases of the proposed CIT, CP-VTON, and CP-VTON+.

the user study according to the same questions asked in Table 1 with randomly selected 50 image sets: (Q1) Which image is the most photo-realistic? (Q2) Which image preserves better the details of the target clothing? And the results shown in Table 3 shows that B3 can generate more realistic images and preserves more details of the clothing items compared to B4. We also marked out the obvious artifacts region in Figure 6.

4.4 Failure Cases and Analysis

Though impressive person try-on images can be generated by our CIT, there are still three kinds of inevitable common failure cases. We visualize them in Figure 7 with a comparison to both CP-VTON and CP-VTON+.

The first case (See the 1st row of Figure 7) is that the difference between the clothing item in the reference image and the in-shop cloth is too large. Consequently, the mask of the person cannot well match the target in-shop clothing item. The second failure case comes from the self-occlusion problem, which leads to blurry ambiguity-prone generated images (See the 2nd row of Figure 7). The third failure case (See the 3rd row of Figure 7) derives from the drastic difference between the pose of the person and the sides of the in-shop cloth. This also leads to ambiguous results. In the first two cases, the main reason may be that the input data lack information in terms of whether the region of a human body should be covered with cloth or not. We propose to give further organization to the input data to remedy this issue, such as using more accurate segmentation maps or adopting more fine-grained human annotations. For the last case, the 2D image-based method cannot completely capture such a complicated relationship between a person and a clothing item. We think that taking the 3D input data, such as body mesh, and 3D clothing items, into consideration may alleviate such a problem.

5 CONCLUSION

In this article, we propose a novel two-stage CIT method for the 2D image-based virtual try-on task. In the first stage, we introduce an interactive Transformer matching block, which is able

to accurately model the global long-range correlations when warping a cloth through the thin-plate spline transformation. Consequently, the quality of the warped clothing item can be more realistic in terms of texture. We also present a transformer-based reasoning block in the second stage for modeling the mutual interactive relations, which can be utilized to further improve the rendering process, resulting in more realistic try-on results. Extensive experiments in terms of quantitative and qualitative comparisons validate that the proposed CIT achieves new competitive performance.

REFERENCES

- [1] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. 2018. Video based reconstruction of 3d people models. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*. 8387–8397.
- [2] Shuai Bai, Huiling Zhou, Zhikang Li, Chang Zhou, and Hongxia Yang. 2022. Single stage virtual try-on via deformable attention flows. In *Proceedings of the European Conference on Computer Vision*. Springer, 409–425.
- [3] Serge Belongie, Jitendra Malik, and Jan Puzicha. 2002. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24, 4 (2002), 509–522.
- [4] Fred L. Bookstein. 1989. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11, 6 (1989), 567–585.
- [5] Remi Brouet, Alla Sheffer, Laurence Boissieux, and Marie-Paule Cani. 2012. Design preserving garment transfer. *ACM Transactions on Graphics* 31, 4 (2012), Article–No.
- [6] Wenzheng Chen, Huan Wang, Yangyan Li, Hao Su, Zhenhua Wang, Changhe Tu, Dani Lischinski, Daniel Cohen-Or, and Baoquan Chen. 2016. Synthesizing training images for boosting human 3d pose estimation. In *Proceedings of the International Conference on 3D Vision*. IEEE, 479–488.
- [7] Seunghwan Choi, Sunghyun Park, Minsoo Lee, and Jaegul Choo. 2021. Viton-hd: High-resolution virtual try-on via misalignment-aware normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14131–14140.
- [8] Yunje Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*. 8789–8797.
- [9] Ayush Chopra, Rishabh Jain, Mayur Hemani, and Balaji Krishnamurthy. 2021. Zflow: Gated appearance flow-based virtual try-on with 3d priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5433–5442.
- [10] Aiyu Cui, Daniel McKee, and Svetlana Lazebnik. 2021. Dressing in order: Recurrent person image generation for pose transfer, virtual try-on and outfit editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14638–14647.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics*.
- [12] Haoye Dong, Xiaodan Liang, Xiaohui Shen, Bowen Wu, Bing-Cheng Chen, and Jian Yin. 2019. Fw-gan: Flow-navigated warping gan for video virtual try-on. In *Proceedings of the International Conference on Computer Vision*. 1161–1170.
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, and others. 2020. An Image is Worth 16x16 Words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.
- [14] Jun Ehara and Hideo Saito. 2006. Texture overlay for virtual clothing based on PCA of silhouettes. In *Proceedings of the IEEE/ACM International Symposium on Mixed and Augmented Reality*. Citeseer, 139–142.
- [15] Benjamin Fele, Ajda Lampe, Peter Peer, and Vitomir Struc. 2022. C-vton: Context-driven image-based virtual try-on network. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 3144–3153.
- [16] Matteo Fincato, Marcella Cornia, Federico Landi, Fabio Cesari, and Rita Cucchiara. 2022. Transform, warp, and dress: A new transformation-guided model for virtual try-on. *ACM Transactions on Multimedia Computing, Communications, and Applications* 18, 2 (2022), 1–24.
- [17] Matteo Fincato, Federico Landi, Marcella Cornia, Fabio Cesari, and Rita Cucchiara. 2021. VITON-GT: An image-based virtual try-on model with geometric transformations. In *Proceedings of the 2020 25th International Conference on Pattern Recognition*. IEEE, 7669–7676.
- [18] Chongjian Ge, Yibing Song, Yuying Ge, Han Yang, Wei Liu, and Ping Luo. 2021. Disentangled cycle consistency for highly-realistic virtual try-on. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16928–16937.

- [19] Xiaoling Gu, Jun Yu, Yongkang Wong, and Mohan S. Kankanhalli. 2020. Toward multi-modal conditioned fashion image translation. *IEEE Transactions on Multimedia* 23, (2020), 2361–2371.
- [20] Peng Guan, Loretta Reiss, David A. Hirshberg, Alexander Weiss, and Michael J. Black. 2012. Drape: Dressing any person. *ACM Transactions on Graphics* 31, 4 (2012), 1–10.
- [21] Erhan Gundogdu, Victor Constantin, Amrollah Seifoddini, Minh Dang, Mathieu Salzmann, and Pascal Fua. 2019. Garnet: A two-stream network for fast and accurate 3d cloth draping. In *Proceedings of the International Conference on Computer Vision*. 8739–8748.
- [22] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S. Davis. 2018. Viton: An image-based virtual try-on network. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*. 7543–7552.
- [23] Sen He, Yi-Zhe Song, and Tao Xiang. 2022. Style-based global appearance flow for virtual try-on. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3470–3479.
- [24] Thibaut Issenhuth, Jérémie Mary, and Clément Calauzenes. 2020. Do not mask what you do not need to mask: A parser-free virtual try-on. In *Proceedings of the Computer Vision—ECCV 2020: 16th European Conference, Glasgow, 23–28, 2020*. Springer, 619–635.
- [25] Paul Jaccard. 1912. The distribution of the flora in the alpine zone. 1. *New Phytologist* 11, 2 (1912), 37–50.
- [26] Nikolay Jetchev and Urs Bergmann. 2017. The conditional analogy gan: Swapping fashion articles on people images. In *Proceedings of the International Conference on Computer Vision Workshops*. 2287–2292.
- [27] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *Proceedings of the European Conference on Computer Vision*. Springer, 694–711.
- [28] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. 2022. Transformers in vision: A survey. *ACM Computing Surveys* 54, 10s (2022), 1–41.
- [29] Zorah Lahner, Daniel Cremers, and Tony Tung. 2018. Deepwrinkles: Accurate and realistic clothing modeling. In *Proceedings of the European Conference on Computer Vision*. 667–684.
- [30] Sangyun Lee, Gyojung Gu, Sunghyun Park, Seunghwan Choi, and Jaegul Choo. 2022. High-resolution virtual try-on with misalignment and occlusion-handled conditions. In *Proceedings of the European Conference on Computer Vision*. Springer, 204–219.
- [31] Kedan Li, Min Jin Chong, Jeffrey Zhang, and Jingen Liu. 2021. Toward accurate and realistic outfits visualization with attention to details. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15546–15555.
- [32] Yu Liu, Wei Chen, Li Liu, and Michael S. Lew. 2019. Swapgan: A multistage generative approach for person-to-person fashion style transfer. *IEEE Transactions on Multimedia* 21, 9 (2019), 2209–2222.
- [33] Yahui Liu, Bin Ren, Yue Song, Wei Bi, Nicu Sebe, and Wei Wang. 2022. Breaking the chain of gradient leakage in vision transformers. arXiv:2205.12551. Retrieved from <https://arxiv.org/abs/2205.12551>
- [34] David G. Lowe. 2004. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* 60, 2 (2004), 91–110.
- [35] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. 2020. Controllable person image synthesis with attribute-decomposed gan. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*. 5084–5093.
- [36] Krystian Mikolajczyk and Cordelia Schmid. 2002. An affine invariant interest point detector. In *Proceedings of the European Conference on Computer Vision*. 128–142.
- [37] M. R. Minar, T. T. Tuan, H. Ahn, P. Rosin, and Y. K. Lai. 2020. Cp-vton+: Clothing shape and texture preserving image-based virtual try-on. In *Proceedings of the Conference on Computer Vision and Pattern Recognition Workshops*, Vol. 2. 11.
- [38] Davide Morelli, Matteo Fincato, Marcella Cornia, Federico Landi, Fabio Cesari, and Rita Cucchiara. 2022. Dress code: High-resolution multi-category virtual try-on. In *Proceedings of the Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VIII*. 345–362.
- [39] Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*. 48–53.
- [40] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J. Black. 2017. ClothCap: seamless 4D clothing capture and retargeting. *ACM Transactions on Graphics* 36, 4 (2017), 1–15.
- [41] Bin Ren, Yahui Liu, Yue Song, Wei Bi, Rita Cucchiara, Nicu Sebe, and Wei Wang. 2023. Masked jigsaw puzzle: A versatile position embedding for vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20382–20391.
- [42] Bin Ren, Hao Tang, and Nicu Sebe. 2021. Cascaded cross MLP-mixer GANs for cross-view image translation. *British Machine Vision Conference* (2021).

- [43] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. 2017. Convolutional neural network architecture for geometric matching. In *Proceedings of the Computer Vision and Pattern Recognition*. 6148–6157.
- [44] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training GANs. In *Proceedings of the Neural Information Processing Systems*.
- [45] Jo Schlemper, Ozan Oktay, Michiel Schaap, Mattias Heinrich, Bernhard Kainz, Ben Glocker, and Daniel Rueckert. 2019. Attention gated networks: Learning to leverage salient regions in medical images. *Medical Image Analysis* 53 (2019), 197–207.
- [46] Cordelia Schmid and Roger Mohr. 1997. Local grayvalue invariants for image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19, 5 (1997), 530–535.
- [47] Yoones A. Sekhavat. 2016. Privacy preserving cloth try-on using mobile augmented reality. *IEEE Transactions on Multimedia* 19, 5 (2016), 1041–1049.
- [48] Masahiro Sekine, Kaoru Sugita, Frank Perbet, Björn Stenger, and Masashi Nishiyama. 2014. Virtual fitting by single-shot body shape estimation. In *Proceedings of the Int. Conf. on 3D Body Scanning Technologies*. Citeseer, 406–413.
- [49] Emma Strubell, Patrick Verga, Daniel Andor, David Weiss, and Andrew McCallum. 2018. Linguistically-informed self-attention for semantic role labeling. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 5027–5038.
- [50] Hao Tang, Song Bai, Philip H. S. Torr, and Nicu Sebe. 2020. Bipartite graph reasoning GANs for person image generation. In *Proceedings of the British Machine Vision Conference*.
- [51] Hao Tang, Song Bai, Li Zhang, Philip H. S. Torr, and Nicu Sebe. 2020. Xinggan for person image generation. In *Proceedings of the European Conference on Computer Vision*. Springer, 717–734.
- [52] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. 2021. MLP-Mixer: An all-MLP architecture for vision. *Advances in Neural Information Processing Systems* (2021), 24261–24272.
- [53] Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, J. Zico Kolter, Louis-Philippe Morency, and Ruslan Salakhutdinov. 2019. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the Neural Information Processing Systems*.
- [55] Bochao Wang, Huabin Zheng, Xiaodan Liang, Yimin Chen, Liang Lin, and Meng Yang. 2018. Toward characteristic-preserving image-based virtual try-on network. In *Proceedings of the European Conference on Computer Vision*. 589–604.
- [56] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. 2018. Non-local neural networks. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*. 7794–7803.
- [57] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612.
- [58] Jun Xu, Yuanyuan Pu, Rencan Nie, Dan Xu, Zhengpeng Zhao, and Wenhua Qian. 2021. Virtual try-on network with attribute transformation and local rendering. *IEEE Transactions on Multimedia* 23, (2021), 2222–2234.
- [59] Han Yang, Ruimao Zhang, Xiaobao Guo, Wei Liu, Wangmeng Zuo, and Ping Luo. 2020. Towards photo-realistic virtual try-on by adaptively generating-preserving image content. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*. 7850–7859.
- [60] Ruiyun Yu, Xiaoqi Wang, and Xiaohui Xie. 2019. Vtnfp: An image-based virtual try-on network with body and clothing feature preservation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10511–10520.
- [61] Miaolong Yuan, Ishtiaq Rasool Khan, Farzam Farbiz, Susu Yao, Arthur Niswar, and Min-Hui Foo. 2013. A mixed reality virtual clothes try-on system. *IEEE Transactions on Multimedia* 15, 8 (2013), 1958–1968.
- [62] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*. 586–595.
- [63] Shizhan Zhu, Raquel Urtasun, Sanja Fidler, Dahua Lin, and Chen Change Loy. 2017. Be your own prada: Fashion synthesis with structural coherence. In *Proceedings of the International Conference on Computer Vision*. 1680–1688.

Received 11 December 2022; revised 28 May 2023; accepted 18 August 2023