

3D Human Texture Estimation from a Single Image with Transformers

Xiangyu Xu Chen Change Loy
S-Lab, Nanyang Technological University
xiangyu.xu@ntu.edu.sg, ccloy@ntu.edu.sg

Abstract

We propose a Transformer-based framework for 3D human texture estimation from a single image. The proposed Transformer is able to effectively exploit the global information of the input image, overcoming the limitations of existing methods that are solely based on convolutional neural networks. In addition, we also propose a mask-fusion strategy to combine the advantages of the RGB-based and texture-flow-based models. We further introduce a part-style loss to help reconstruct high-fidelity colors without introducing unpleasant artifacts. Extensive experiments demonstrate the effectiveness of the proposed method against state-of-the-art 3D human texture estimation approaches both quantitatively and qualitatively. The project page is at <https://www.mmlab-ntu.com/project/textformer>.

1. Introduction

In this paper, we study the problem of estimating 3D human texture from a monocular image. This is an important problem that plays a key role in single-image 3D human reconstruction and has wide applications in virtual and augmented reality, film industry, gaming, and biometrics.

Most of existing methods for this problem [13, 23, 29, 44, 47, 53] use deep convolution neural networks (CNNs) to predict 3D human texture (*i.e.*, a UV map) from the input image (Figure 1(a)). While these methods have achieved impressive results, their network architectures suffer from an inherent shortcoming: convolution layers are by design local operations and inefficient in processing global information that is crucial in 3D human texture estimation. More specifically, the input and output in this task do not have strictly-aligned spatial correspondences and may even have totally different shapes as shown in Figure 1(a). This is in sharp contrast to 2D computer vision tasks such as image super-resolution [51] and image-to-image translation [21] where the input and output are well aligned.

We believe that one should not simply follow the common practice of using only CNNs for the 3D task in this

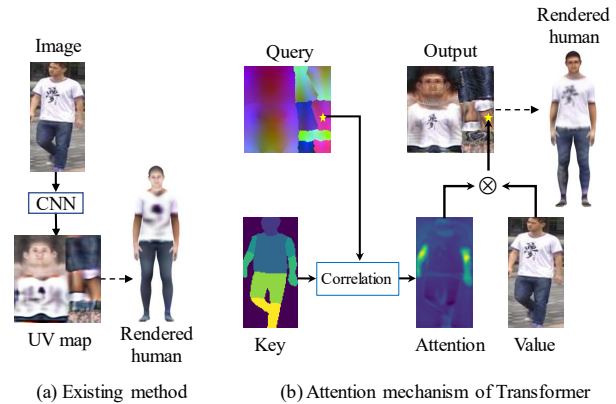


Figure 1. We propose a Transformer for 3D human texture estimation from a single image. Compared with the existing method [44] that solely relies on CNNs (a), the proposed Transformer uses the attention mechanism to more effectively exploit global information (b), which leads to higher-quality 3D human texture estimation. See the text for more detailed explanations.

work, and global operations should be introduced for better 3D human texture reconstruction. Note that some recent approaches [47, 53] attempt to address this issue with fully-connected layers (or MLPs), which however often leads to loss of fine spatial information. Moreover, this information loss cannot be easily remedied by skip connections [38] due to the misalignment of the different-layer features. A successful solution should be able to more effectively distribute features of the input into suitable locations in the UV space while preserving fine spatial information.

Towards this end, we propose a Transformer-based framework for 3D human texture estimation from a single image. The Transformer allows processing information of the input in a global manner, which is in particular suitable for our task. At the core of the Transformer is an attention module that involves three basic components: Query, Key, and Value. For the Query, we use a pre-computed color map that has the same shape as the output UV map. Each pixel in the query map corresponds to a vertex on 3D human mesh [31]. The Value is the input image that has all the source pixels. For the Key, we use a 2D part-segmentation map obtained with an off-the-shelf model [16], which implies the mapping from the image to the UV space [34, 53].

For an intuitive understanding of the relationship between these three components, we elucidate the working mechanism of our Transformer in Figure 1(b). A pixel in the Query (marked as yellow star) is first used to correlate with the Key, which produces an attention map. With this attention map, the source information in the Value can then be effectively aggregated by weighted averaging to generate the corresponding pixel in the output UV map. Such an attention mechanism allows us to exploit global information of the input image without losing fine details, which is the key factor that distinguishes our method from existing algorithms. Note that the above explanations have been simplified for ease of understanding. As will be introduced in Section 3, we use more channels for the Value and Key in real implementation and perform multi-head attention [43] in feature space.

To summarize, we make the following contributions:

- 1) We propose a Transformer-based framework, termed as Texformer, for 3D human texture estimation from a single image. Based on the attention mechanism, the proposed network is able to effectively exploit global information of the input. It naturally overcomes the limitations of existing algorithms that solely rely on CNNs and effectively facilitates higher-quality 3D human texture reconstruction.
- 2) Existing algorithms output either RGB values [44, 47] or texture flow [23, 53] to synthesize the final UV map. We analyze the limitations of these two strategies and propose a new method to combine the best of both worlds. We show that the proposed method is able to significantly reduce visual artifacts while preserving fine details.
- 3) The estimated textures of previous approaches often suffer from noticeable color differences from the input. To remedy this issue, we propose a part-style loss that enforces the Gram-matrix similarity [12] for each human body part and encourages closer appearances to the input image.

Extensive experiments on the Market-1501 dataset [54] demonstrate the effectiveness of the proposed method against state-of-the-art 3D human texture estimation approaches both quantitatively and qualitatively.

2. Related Work

3D human texture estimation. Recent years have witnessed significant progress in the field of 3D human texture estimation [1, 2, 3, 6, 13, 19, 23, 28, 29, 30, 34, 35, 36, 39, 44, 47, 53, 55, 57]. Some methods [1, 2, 3, 6, 34] solve this problem by taking multi-view images as input, and the 3D human texture can be synthesized by merging textures from different views with Graph-cut [7], median filters [3], or deep learning [34]. However, these methods cannot be easily used for single-image 3D human reconstruction [4, 37, 40, 56, 59]. More recent approaches in this direction use RGBD videos [30, 57], which poses even higher demand for the input.

For applications in single-image 3D human reconstruction [4, 37, 40, 56, 59], another line of approaches use deep neural networks to reconstruct 3D human textures from a monocular image [19, 28, 35, 39, 55]. However, these methods usually need 3D supervision obtained via 3D scanning, which is both time-consuming and labor-intensive. Moreover, some methods [28, 36] require high-quality dense human pose estimation for extracting partially observed textures, which would be challenging for in-the-wild images.

More closely related to this work, some recent methods aim to address the problem of single-image texture reconstruction without 3D labels [13, 23, 29, 44, 47, 53], which are generally more applicable for real-world usage. For instance, Wang *et al.* [44] propose a self-supervised training pipeline by building upon the success of existing 3D human mesh reconstruction algorithms [18, 22, 24, 26, 27, 46]. Kanazawa *et al.* [23] also learn the 3D textures in a self-supervised manner and use texture flow to directly sample pixels from the input image. In addition, Zhao *et al.* [53] exploit human part segmentation as input and use multi-view data as extra training labels for better performance. While these methods have achieved promising progress in this task, they usually rely on CNNs that cannot well exploit global information of the input. In contrast, we propose a Transformer-based framework that effectively overcomes this issue. In addition, we introduce a mask-fusion strategy to combine the advantages of both RGB predictions and texture flow. We also propose part-style loss and face-structure loss to further improve the texture reconstruction results.

Transformers. Thanks to the remarkable ability to handle long-range information, the Transformer has become the dominant architecture in natural language processing [10, 43]. Recently, it has also been introduced in computer vision, and the applications include object detection [8], image classification [11], image restoration [9], and hand pose estimation [17]. We refer to [14] for a more comprehensive review of vision transformers.

In this work, we demonstrate that the attention mechanism makes the Transformer particularly suitable for the task of 3D human texture estimation. We provide new insights in designing the Transformer for this task, and the proposed network achieves better performance than the state-of-the-art approaches. We also show that a common Transformer structure without specific design cannot perform as well as the proposed architecture.

3. Methodology

We propose a Transformer-based framework, termed as Texformer, for 3D human texture estimation from a single image. An overview of the Texformer is shown in Figure 2. A brief explanation of the framework is provided in the figure caption. Next, we provide more detailed explanations of the Texformer as well as our loss functions.

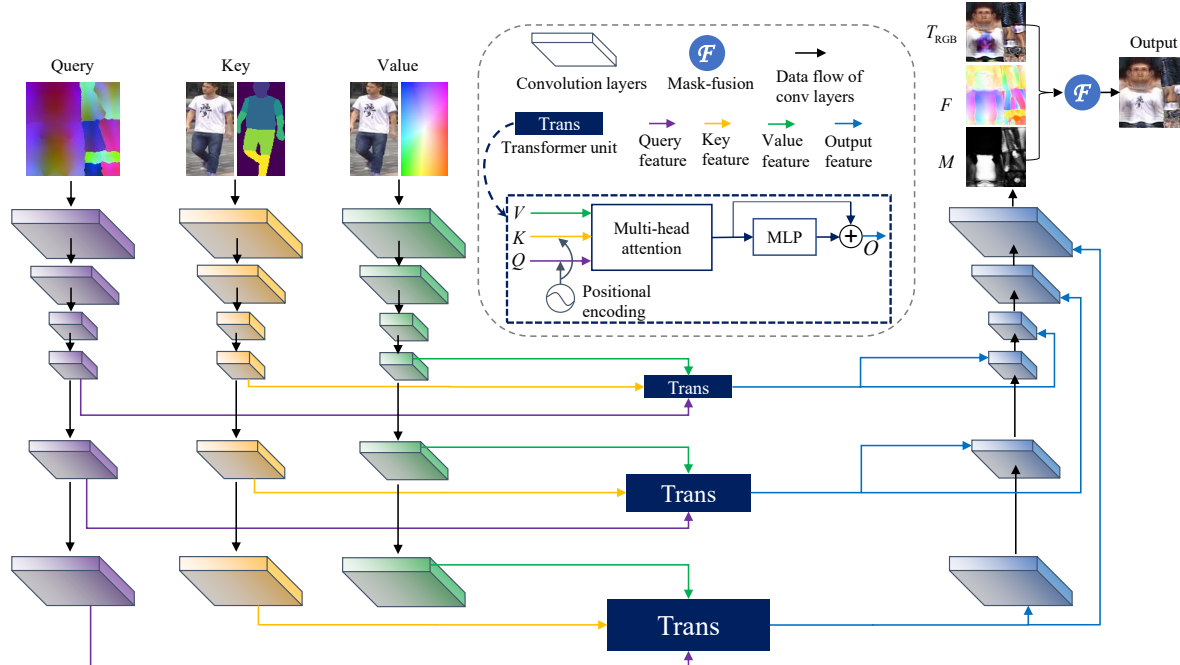


Figure 2. Overview of the proposed Texformer. The Query is a pre-computed color encoding of the UV space obtained by mapping the 3D coordinates of a standard human body mesh to the UV space. The Key is a concatenation of the input image and the 2D part-segmentation map. The Value is a concatenation of the input image and its 2D coordinates. We first feed the Query, Key, and Value into three CNNs to transform them into feature space. Then the multi-scale features are sent to the Transformer units to generate the Output features. The multi-scale Output features are processed and fused in another CNN, which produces the RGB UV map T_{RGB} , texture flow F , and fusion mask M . The final UV map is generated by combining T_{RGB} and the textures sampled with F using the fusion mask M . Note that we have skip connections between the same-resolution layers of the CNNs similar to [48] which have been omitted in the figure for brevity.

3.1. Texformer

Query. As introduced in Section 1, the Query of our model represents a color-encoding of the output UV space, and each pixel in the Query map should characterize a vertex on the 3D human mesh. To this end, we use the 3D coordinates of a standard human body mesh from the SMPL model [31] as the color encoding of each vertex. The color values of each pixel in the Query can be obtained by first mapping the 3D coordinates to the UV space and then interpolating the coordinates with KD-Tree [33]. We pre-compute the Query map before network training, and the same color encoding of the UV space is used for all input images.

Key. Recall that the Key is used to correlate with the Query elements to obtain the attention map for the input (Figure 1), which is crucial for connecting the image space and the UV space. Motivated by recent studies [34, 53], we use 2D part-segmentation as the Key to learn a mapping from the input image to the output UV map. Our 2D part-segmentation is obtained with an off-the-shelf model [16]. As shown in Figure 2, we also include the input image as additional channels of the Key to provide contextual information.

Value. The Value represents the source information indexed by the Key and is aggregated into the UV space to generate the output using the attention map. There are two possible

options for the Value depending on the form of the Transformer output: 1) using the RGB input image as the Value when the Transformer directly outputs the RGB UV map; 2) using the flow field of the image, *i.e.*, 2D coordinates for each pixel, when the model first produces the texture flow and then generates the UV map by sampling the input image with the predicted flow. As will be introduced later in this section, our Transformer predicts both RGB values and texture flow, and thus, we concatenate both the RGB image and the flow field as the Value of our model (Figure 2).

The Query, Key, and Value cannot be directly used by the Transformer with their raw forms; this is especially the case for the Query and Key which will be compared in the attention module and thereby should have the same feature dimension. Thus, we feed them in CNNs to transform them into feature space as shown in Figure 2. Then the resulting features $Q \in \mathbb{R}^{vu \times d}$, $K \in \mathbb{R}^{hw \times d}$, and $V \in \mathbb{R}^{hw \times c}$ are sent into the Transformer unit, which produces the Output features $O \in \mathbb{R}^{vu \times c}$. Here, v, u are the height and width of the output UV map, h, w are the height and width of the input image, while d, c are the feature dimensions of K and V , respectively. Note that the CNN of Q is used to find a better representation of the fixed Query map and only needed in the training phase. During deployment, we can drop the Query CNN by pre-computing the feature encodings.

The Transformer unit in Figure 2 is the central part of the proposed network, which effectively distributes the image features into suitable locations of the UV map and enables global information interchange between the input space and the output space. Specifically, it can be written as:

$$O = f_{\text{res-MLP}}(f_{\text{Attn}}(Q + E_Q, K + E_K, V)), \quad (1)$$

where f_{Attn} is the multi-head attention module from [43]. $f_{\text{res-MLP}}$ is a two-layer MLP with a residual connection between its input and output. E_Q and E_K are sinusoidal positional encodings [43] for the Query and Key features, respectively.

Low-rank attention. As introduced in [43], the multi-head attention f_{Attn} is based on a single-head attention module:

$$f_{\text{single}}(\tilde{Q}, \tilde{K}, \tilde{V}) = \text{softmax}(\tilde{Q}\tilde{K}^\top / \alpha)\tilde{V}, \quad (2)$$

where $\tilde{Q} \in \mathbb{R}^{\tilde{v}\tilde{u} \times \tilde{d}}$, $\tilde{K} \in \mathbb{R}^{\tilde{h}\tilde{w} \times \tilde{d}}$, and $\tilde{V} \in \mathbb{R}^{\tilde{h}\tilde{w} \times \tilde{c}}$. The denominator α is incorporated to avoid large values. For clarity, we do not use the symbols Q , K , and V here as these features need to be first processed by linear projection layers in the multi-head attention f_{Attn} [43].

One important problem of the normal attention formulation is the high memory complexity ($\mathcal{O}(\tilde{v}\tilde{u}\tilde{h}\tilde{w})$) for Eq. 2, which makes the training infeasible for common GPUs. To remedy this issue, we propose a low-rank attention (LoRA) module inspired by [50]:

$$f_{\text{LoRA}}(\tilde{Q}, \tilde{K}, \tilde{V}) = \tilde{Q}\tilde{K}^\top\tilde{V} / \alpha, \quad (3)$$

where the softmax function is removed, and the attention is approximated by pure matrix multiplication. This strategy allows more efficient computations by manipulating the order of the matrix multiplications [50] (*i.e.*, computing $\tilde{K}^\top\tilde{V}$ first). The memory footprint can be significantly reduced to $\mathcal{O}(\max(\tilde{v}\tilde{u}, \tilde{h}\tilde{w}) \cdot \max(\tilde{d}, \tilde{c}))$, which makes the training more feasible. Note that [50] uses this low-rank strategy for matrix factorization while we apply it to multi-head attention for efficient Transformers.

Multi-scale feature fusion. Instead of applying the Transformer unit at one single scale, we compute the Output features with a three-level feature pyramid. For the first two levels where the features have large spatial sizes, we use the proposed LoRA (Eq. 3) to reduce memory cost. For the small features on the third level, we simply use the softmax attention module (Eq. 2). Then the multi-scale Output features $O_i, i = 1, 2, 3$ can be easily fused within a CNN as shown in Figure 2. This multi-scale approach can help the proposed Transformer better exploit contextual information and predict higher-quality human textures for invisible regions (Figure 9).

Mask-fusion of RGB and texture flow. There are two possible options for the output of the network: 1) directly generating the RGB textures T_{RGB} [44, 47], or 2) predicting the

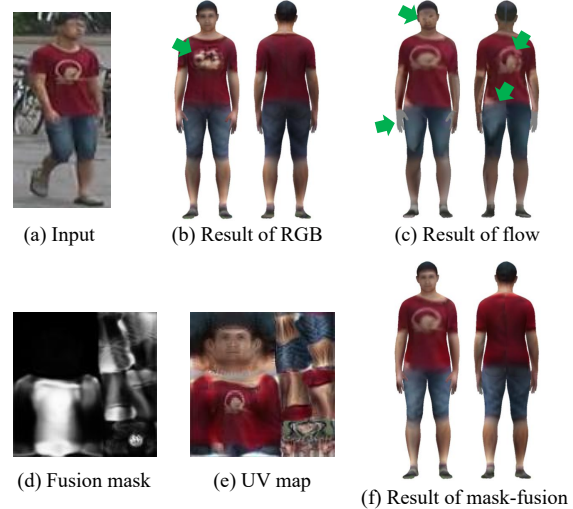


Figure 3. Mask-fusion strategy combines RGB values and texture flow to reduce artifacts while preserving fine details.

texture flow F that can be used to generate the final textures by sampling from the input image I [23, 49, 53]. As shown in Figure 3, both these two strategies have their advantages and limitations: Directly synthesizing T_{RGB} can well reconstruct visually-pleasant 3D human textures, which however often leads to loss of fine details (Figure 3(b)). On the other hand, the texture flow is able to preserve fine details in the input image, but the results could suffer from severe artifacts (Figure 3(c)).

To solve this problem, we propose a mask-fusion method to combine the advantages of both RGB values and texture flow. Specifically, our Transformer generates three outputs: the RGB UV map T_{RGB} , the texture flow F , and a fusion mask M . Then the mask-fusion process can be written as:

$$T = M \odot f_{\text{sample}}(F, I) + (1 - M) \odot T_{\text{RGB}}, \quad (4)$$

where f_{sample} is the bilinear sampling function that is used to sample textures from I with F , and \odot denotes element-wise multiplication. The final texture T is computed by fusing T_{RGB} and $f_{\text{sample}}(F, I)$ via weighted sum. As shown in Figure 3(f), the mask-fusion method is able to significantly reduce visual artifacts while preserving fine details. In addition, by observing the learned fusion mask (Figure 3(d)), we can find that the network tends to use sampled textures $f_{\text{sample}}(F, I)$ for visible regions such as the human chest, and use synthesized textures T_{RGB} for invisible regions such as the human back as well as other hard regions with sophisticated structures such as the face and hands.

Relationship with conventional Transformers. Typically, existing works [8, 9, 11, 43] construct the Transformer by stacking multiple Transformer units, where the Output features of the current unit become the Query features for the next one. Thus, the data flow in the Transformer needs to be

able to serve as two different roles (*i.e.*, Query and Output) simultaneously. This is a reasonable choice when either 1) the Query and Output are in the same space, for example, in natural language processing [43] the Query is the shifted Output of the previous time step, or 2) the physical meaning of the Query cannot be concretely defined, for instance, in object detection [8] the Query is randomly-initialized vectors that are also optimized during training.

However, in this work we have a clear definition of the Query which is a color encoding map obtained from 3D coordinates of a standard 3D body mesh, and each element in the Query corresponds to a physical vertex; and this Query map is in a space different from the Output (*i.e.*, the RGB values and texture flow). Therefore, in our network (Figure 2), we avoid stacking multiple attention modules and use the Transformer units more efficiently in a multi-scale manner. Through this way, the features of our Texformer do not need to serve as two different roles simultaneously anymore, which effectively alleviates the difficulties of training. As shown in Section 4.1, the conventional Transformer cannot provide as high-quality results as the proposed model, which demonstrates the effectiveness of our design.

3.2. Loss Functions

Following [23, 44, 47, 53], we train the Texformer in a self-supervised manner. First, we use the estimated human textures to render the human image $f_r(T, D)$, where f_r is a differentiable rendering function [25], and D represents the 3D human mesh and camera parameters predicted by the state-of-the-art algorithm RSC-Net [46]; then the model can be trained by enforcing the similarity between the rendered image $f_r(T)$ and the input I . Note that we omit D here and in the following sections for conciseness.

To enforce the similarity between $f_r(T)$ and I , we first use the re-identification (ReID) loss from [44]:

$$\ell_{\text{ReID}} = \sum_v \|\bar{\phi}_v(f_r(T)) - \bar{\phi}_v(I)\|_2^2, \quad (5)$$

where ϕ_v represents the v -th layer of a pedestrian re-identification network [41]. $\bar{\phi} = \phi/\|\phi\|_2$ denotes L2-normalization.

While this ReID loss performs reasonably well, the reconstructed textures often have noticeable color differences from the input image as shown in Figure 4(a), which is mainly due to the normalization operation in Eq. 5. One way to solve this problem is to remove the normalization and use the unnormalized features for the ReID loss. However, this strategy often leads to severe artifacts in the reconstructed human textures as shown in Figure 4(b) (note the unnatural textures on human face and arms).

We suppose this problem could be owing to that the predicted 3D human mesh is not always accurate. As shown in




Figure 4. Part-style loss for reconstructing human textures with high-fidelity colors.




Figure 5. Misalignment issue of the 3D mesh predictions. (b) is the 3D body mesh predicted by [46]. (c) is the 2D body part segmentation (P-S) obtained from the 3D mesh (b). (d) is the 2D body part segmentation obtained by [16].

Figure 5(b), the 3D human mesh does not align well with the input image. Therefore, directly enforcing the similarity between the rendered human and the input image leads to significant errors, which will negatively affect the training process.

Part-style loss. To address this issue, we propose a new loss function based on the observation that the 2D part-segmentation is generally more accurate than 3D estimation and aligns better with the input image (Figure 5(d)). Specifically, we enforce the similarity between each body part of the rendered human and the input image, which naturally handles the misalignment issue caused by inaccurate 3D mesh predictions.

However, the regions of the same body part in different segmentation maps usually have different sizes and shapes as shown in Figure 5(c) and (d), which prevents us from using a simple MSE loss as in Eq. 5. Instead, we employ the style loss from [12] where the Gram matrix does not require the same size and shape for computation. Different from the original style loss [12] that computes the Gram matrix for the whole image, we enforce the Gram-matrix similarity in a per-body-part manner, which can be formulated as:

$$\ell_{\text{style}} = \sum_p \|G(M_p \odot \phi_1(f_r(T))) - G(M'_p \odot \phi_1(I))\|_2^2,$$

where M and M' are the human part segmentation from the 3D mesh and the 2D human parsing model [16], respectively. Here, p indicates the p -th body part, and G is the Gram matrix. Note that we only use the features from the

first layer of the ReID network (*i.e.*, ϕ_1) for the part-style loss to better focus on low-level colors of the reconstructed textures. As shown in Figure 4(c), combining ℓ_{ReID} and ℓ_{style} achieves color appearances closer to the input image without introducing unpleasant artifacts.

Face-structure loss. While our method is able to well reconstruct the 3D textures for most regions of the human body, it is still challenging to synthesize the textures for the human parts with complicated structures, such as the human face. Wang *et al.* [44] propose an MSE face loss to encourage the face textures to be close to the mean face of a synthetic human dataset [42], which however, often leads to unnatural results where the face color is not consistent with other skin regions of the reconstructed human. To remedy this issue, we propose a face-structure loss:

$$\ell_{\text{face}} = -\frac{1}{N} \sum_{i=1}^N s(M_{\text{face}} \odot T, M_{\text{face}} \odot T_{\text{syn}}^{(i)}), \quad (6)$$

where $\{T_{\text{syn}}^{(i)}\}_{i=1}^N$ is a collection of realistic human textures from a synthetic human dataset [42]. M_{face} is a pre-defined binary map that indicates the face region on the texture map. s is a structure-similarity function [45] defined as:

$$s(x, y) = (\sigma_{xy} + C) / (\sigma_x \sigma_y + C), \quad (7)$$

where σ_x is the standard deviation of x , σ_{xy} is the covariance between x and y , and C is a constant. Eq. 6 essentially encourages the face textures to have similar structures as T_{syn} , which effectively facilitates the generation of plausible face textures while retaining the colors of the input human.

Our final loss is a combination of the ReID loss, part-style loss, and face-structure loss:

$$\ell = w_1 * \ell_{\text{ReID}} + w_2 * \ell_{\text{style}} + w_3 * \ell_{\text{face}}, \quad (8)$$

where w_1 , w_2 , and w_3 are hyper-parameters. Similar to [53], we use multi-view data to train the network, and our loss functions can be straightforwardly applied to multi-view images.

4. Experiments

We use the Market-1501 dataset [54] in our experiments, which consists of human images of 1501 person identities. Following [44, 47], we use 1401 identities for training and the remaining 100 identities for testing. For the face-structure loss, we use synthetic textures from the SUR-REAL dataset [42]. We use 8 attention heads for the attention modules, and the feature dimension is set as 128 for both K and V . We set $\alpha = \hat{h}\hat{w}$ for the LoRA (Eq. 3). We use BatchNorm [20] instead of LayerNorm [5] for the Transformer units as it gives better results in practice. Similar to [44, 47], we use the PCB network [41] as the feature

Table 1. Quantitative evaluation of the proposed algorithm. Numbers in bold indicate the best in each column, and underscored numbers indicate the second.

Method	CosSim \uparrow	CosSim-R \uparrow	SSIM \uparrow	LPIPS \downarrow	Params (M)
CMR [23]	0.5241	0.4978	0.7142	0.1275	84.1
HPBTT [53]	0.5246	0.5027	<u>0.7420</u>	<u>0.1168</u>	235.0
RSTG [44]	0.5282	0.4924	0.6735	0.1778	13.4
TexGlo [47]	0.5408	0.5048	0.6658	0.1776	16.1
Texformer	0.5747	0.5422	0.7422	0.1154	7.6
DETR [8]	0.4344	0.4064	0.6482	0.2409	41.6
DETR* [8]	<u>0.5632</u>	<u>0.5274</u>	0.7133	0.1379	17.9

extractor of the ReID loss (Eq. 5). We empirically set the coefficients in Eq. 8 as $w_1 = 5000$, $w_2 = 0.4$, $w_3 = 0.01$.

For the evaluation metrics, we use the SSIM [45] and LPIPS [52] with human masks [32] to measure the quality of the predicted human textures. We also compute the cosine similarity (CosSim) of person ReID features [41] to measure the results from a semantic level. A higher CosSim indicates that the rendered human is more likely to be the same person in the input image. As the PCB network [41] has been used for training the baselines and our model, we also compute the cosine similarity of the features from a ReID network unseen during training. Specifically, we use the ResNet-50 [15] from [58] and name the corresponding metric as CosSim-R.

4.1. Comparison with the State of the Arts

We compare against the state-of-the-art 3D texture estimation methods: CMR [23], HPBTT [53], RSTG [44], and TexGlo [47]. As HPBTT splits data in a different way, we retrain it with our protocol following [44, 47]. As shown in Table 1, the proposed method achieves consistently better results than the baseline approaches on all metrics while requiring a smaller amount of parameters, which demonstrates the effectiveness of our algorithm.

In addition, we also provide qualitative evaluations of our method against the baselines in Figure 6. As the textures of CMR [23] and HPBTT [53] are sampled from the input image with texture flow, they usually have colors close to the input, which leads to good performance in terms of SSIM and LPIPS in Table 1. However, these methods often suffer from significant artifacts and are not robust to occlusions as shown in Figure 6(a)-(b). On the other hand, while RSTG [44] (Figure 6(c)) and TexGlo [47] (Figure 6(d)) can well reconstruct human textures without severe visual artifacts, their results usually lack fine details and have significant color differences from the input. In contrast, the proposed algorithm achieves higher-quality results with fine details and high-fidelity colors in Figure 6(e).

User study. We conduct a user study for a more comprehensive evaluation of the texture estimation algorithms. This study uses 10 images randomly selected from the Market-1501 test set [54], and each input is processed by 5 different methods: CMR [23], HPBTT [53], RSTG [44], Tex-




Figure 6. Qualitative evaluation of the proposed algorithm.

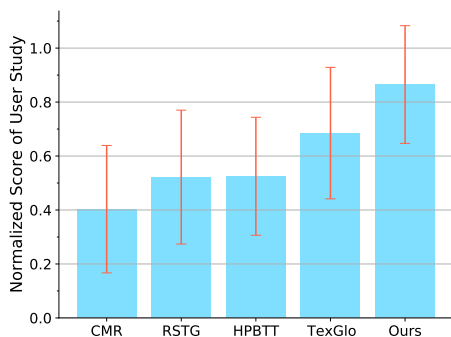


Figure 7. User study of different algorithms for 3D human texture estimation. Each bar shows the mean and standard deviation of the normalized scores of each method.

Glo [47], and the proposed Texformer. Twenty subjects are asked to rank the reconstructed textures by each method with the input image as reference (1 for poor and 5 for excellent). We normalize the rank values to [0, 1], and use the normalized scores for measuring the results. We visualize the mean scores and the standard deviation in Figure 7. The proposed method is clearly preferred over other methods, suggesting its better capability in generating 3D human textures with high perceptual quality.

Comparison with DETR. The DETR [8], which employs a conventional Transformer structure [43], is originally used for object detection. Here we adapt this network for 3D human texture estimation by replacing the object queries with UV map queries, which are randomly initialized and then optimized during training. We also add the proposed Query




Figure 8. Visual comparison with DETR [8].

map to the queries of DETR for better performance. However, the original network cannot work properly as shown in Table 1 and Figure 8(a), which is mainly due to the very deep structure (ResNet-50 [15]) of DETR that prevents the preservation of lower-level texture information. Therefore, we improve the original DETR by only using the first block of the ResNet-50 as the feature extractor. This leads to a stronger baseline model (DETR*) as shown in Table 1 and Figure 8(b). However, as DETR stacks multiple attention modules together without explicitly disentangling the roles of the Query and Output features, it does not perform as well as the proposed network which produces high-quality human textures with fine details as shown in Figure 8(c).

4.2. Ablation Study

Effectiveness of the Transformer. The proposed network essentially relies on the Transformer unit (Eq. 1) to exploit the global information of the input. To analyze its effect, we remove the Transformer units in Figure 2 and instead gen-

Table 2. Ablation study of the proposed method.

Method	CosSim \uparrow	CosSim-R \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o Transformer unit	0.5413	0.5101	0.7242	0.1235
w/o multi-scale	0.5715	0.5391	0.7407	0.1181
only RGB	0.5711	0.5370	0.7167	0.1261
only texture flow	0.5515	0.5258	0.7581	0.1027
w/o part-style loss	0.5783	0.5441	0.7158	0.1412
Full model	<u>0.5747</u>	<u>0.5422</u>	<u>0.7422</u>	<u>0.1154</u>




Figure 9. Effectiveness of different components of Texformer.




Figure 10. Visualization of the learned attention map. The red and yellow stars in the UV map correspond to the attention maps surrounded by red and yellow frames, respectively.

erate the Output features by first concatenating Q , K , and V and then sending the concatenated features into a normal convolution layer. As shown in Table 2, the model without the Transformer unit suffers from a significant performance drop in terms of all the metrics. We also provide a visual example in Figure 9, where the result of “w/o Transformer unit” has much lower visual quality with unpleasant artifacts. To better understand the attention mechanism of the Transformer unit, we further present a visualization of the learned attention map in Figure 10, showing that the Transformer unit can learn to aggregate relevant features from the image into the UV space.

Furthermore, in our approach we use a multi-scale strategy to fuse the features from the Transformer units, which moderately improves the performance as shown in Table 2. As the multi-scale strategy can effectively exploit the contextual information of the input, it is able to better infer textures for invisible regions of the human body, e.g. the human back in Figure 9.

Effectiveness of combining RGB and texture flow. As introduced in Section 3.1, we propose a mask-fusion method to combine the RGB output and texture flow. As shown in Table 2, directly predicting the RGB values leads to poor SSIM and LPIPS, which is mainly due to the incapability of the RGB-based model to reconstruct accurate details (Figure 3(b)). Furthermore, only using texture flow for human




Figure 11. Effectiveness of the part-style loss.

texture estimation results in low CosSim and CosSim-R in Table 2, which is largely caused by the significant amount of artifacts as in Figure 3(c). Note that while both our RGB-based and flow-based models have their limitations, they still achieve better results than the previous RGB-based and flow-based approaches (Table 1), which again shows the superiority of the proposed Transformer network.

Effectiveness of the part-style loss. As shown in Table 2, the part-style loss significantly improves the SSIM and LPIPS. In Figure 11, the reconstructed textures without this loss have noticeable color differences from the input image. This is consistent with the motivation of our design in Section 3.2 to synthesize high-fidelity colors without introducing unpleasant visual artifacts. Note that the part-style loss leads to a slight decrease of the CosSim and CosSim-R in Table 2. This is possibly due to that the ReID network does not fully rely on color appearances to recognize humans, as the same person may appear in different colors when viewed from different angles or under different light conditions. Therefore, enforcing the output to have close color appearances to the input may not always coincide with the optimization direction of the ReID loss, and may slightly distract the training process.

5. Concluding Remarks

We have developed the Texformer for high-quality 3D human texture reconstruction from a single image. At the core of this work is the Transformer unit that allows efficient information interchange between two different spaces, *i.e.* the image space and UV texture space. Potentially, this idea could be extended to other fields where the data of interest involves multiple representations or lies in different spaces, such as multi-modality learning and graph processing.

On the other hand, due to the limitation of SMPL [31], the proposed method is mostly suitable for tight-fitting clothes rather than loose-fitting ones. It is possible to overcome this limitation by introducing more advanced human body models in Texformer, which could be an interesting direction for future research.

Acknowledgement. This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single rgb camera. In *CVPR*, 2019. 2
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *International Conference on 3D Vision (3DV)*, 2018. 2
- [3] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, 2018. 2
- [4] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *ICCV*, 2019. 2
- [5] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 6
- [6] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3D people from images. In *ICCV*, 2019. 2
- [7] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *TPAMI*, 23(11):1222–1239, 2001. 2
- [8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 2, 4, 5, 6, 7
- [9] Hanqing Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. *arXiv preprint arXiv:2012.00364*, 2020. 2, 4
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 4
- [12] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016. 2, 5
- [13] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. In *ECCV*, 2020. 1, 2
- [14] Kai Han, Yunhe Wang, Hanqing Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chunjing Xu, Yixing Xu, et al. A survey on visual transformer. *arXiv preprint arXiv:2012.12556*, 2020. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6, 7
- [16] Houjing Huang, Wenjie Yang, Xiaotang Chen, Xin Zhao, Kaiqi Huang, Jinbin Lin, Guan Huang, and Dalong Du. Eanet: Enhancing alignment for cross-domain person re-identification. *arXiv preprint arXiv:1812.11369*, 2018. 1, 3, 5
- [17] Lin Huang, Jianchao Tan, Ji Liu, and Junsong Yuan. Hand-transformer: Non-autoregressive structured modeling for 3D hand pose estimation. In *ECCV*, 2020. 2
- [18] Zeng Huang, Tianye Li, Weikai Chen, Yajie Zhao, Jun Xing, Chloe LeGendre, Linjie Luo, Chongyang Ma, and Hao Li. Deep volumetric video from very sparse multi-view performance capture. In *ECCV*, 2018. 2
- [19] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *CVPR*, 2020. 2
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 6
- [21] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 1
- [22] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 2
- [23] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018. 1, 2, 4, 5, 6
- [24] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *CVPR*, 2019. 2
- [25] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3D mesh renderer. In *CVPR*, 2018. 5
- [26] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, 2020. 2
- [27] Nikos Kolotouros, Georgios Pavlakos, Michael J. Black, and Kostas Daniilidis. Learning to reconstruct 3D human pose and shape via model-fitting in the loop. In *ICCV*, 2019. 2
- [28] Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 360-degree textures of people in clothing from a single image. In *International Conference on 3D Vision (3DV)*, 2019. 2
- [29] Xueting Li, Sifei Liu, Kihwan Kim, Shalini De Mello, Varun Jampani, Ming-Hsuan Yang, and Jan Kautz. Self-supervised single-view 3D reconstruction via semantic consistency. In *ECCV*, 2020. 1, 2
- [30] Zhe Li, Tao Yu, Chuanyu Pan, Zerong Zheng, and Yebin Liu. Robust 3d self-portraits in seconds. In *CVPR*, 2020. 2
- [31] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics*, 34(6):248, 2015. 1, 3, 8
- [32] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *NeurIPS*, 2017. 6
- [33] Songrit Maneewongvatana and David M Mount. Analysis of approximate nearest neighbor searching with clustered point sets. In *ALLENEX*, 1999. 3
- [34] Aymen Mir, Thiemo Alldieck, and Gerard Pons-Moll. Learning to transfer texture from clothing images to 3D humans. In *CVPR*, 2020. 1, 2, 3
- [35] Ryota Natsume, Shunsuke Saito, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, and Shigeo Morishima. Siclope: Silhouette-based clothed people. In *CVPR*, 2019. 2

- [36] Natalia Neverova, Riza Alp Guler, and Iasonas Kokkinos. Dense pose transfer. In *ECCV*, 2018. [2](#)
- [37] Albert Pumarola, Jordi Sanchez-Riera, Gary Choi, Alberto Sanfeliu, and Francesc Moreno-Noguer. 3DPeople: Modeling the geometry of dressed humans. In *ICCV*, 2019. [2](#)
- [38] 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*, 2015. [1](#)
- [39] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. [2](#)
- [40] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *CVPR*, 2020. [2](#)
- [41] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *ECCV*, 2018. [5](#), [6](#)
- [42] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *CVPR*, 2017. [6](#)
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. [2](#), [4](#), [5](#), [7](#)
- [44] Jian Wang, Yunshan Zhong, Yachun Li, Chi Zhang, and Yichen Wei. Re-identification supervised texture generation. In *CVPR*, 2019. [1](#), [2](#), [4](#), [5](#), [6](#)
- [45] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 13(4):600–612, 2004. [6](#)
- [46] Xiangyu Xu, Hao Chen, Francesc Moreno-Noguer, Laszlo A Jeni, and Fernando De la Torre. 3D human shape and pose from a single low-resolution image with self-supervised learning. In *ECCV*, 2020. [2](#), [5](#)
- [47] Xiangyu Xu, Hao Chen, Francesc Moreno-Noguer, Laszlo A Jeni, and Fernando De la Torre. 3D human pose, shape and texture from low-resolution images and videos. *TPAMI*, 2021. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [48] Xiangyu Xu, Muchen Li, Wenxiu Sun, and Ming-Hsuan Yang. Learning spatial and spatio-temporal pixel aggregations for image and video denoising. *TIP*, 29:7153–7165, 2020. [3](#)
- [49] Xiangyu Xu, Siyao Li, Wenxiu Sun, Qian Yin, and Ming-Hsuan Yang. Quadratic video interpolation. In *NeurIPS*, 2019. [4](#)
- [50] Xiangyu Xu, Yongrui Ma, and Wenxiu Sun. Learning factorized weight matrix for joint filtering. In *ICML*, 2020. [4](#)
- [51] Xiangyu Xu, Deqing Sun, Jinshan Pan, Yujin Zhang, Hanspeter Pfister, and Ming-Hsuan Yang. Learning to super-resolve blurry face and text images. In *ICCV*, 2017. [1](#)
- [52] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [6](#)
- [53] Fang Zhao, Shengcai Liao, Kaihao Zhang, and Ling Shao. Human parsing based texture transfer from single image to 3D human via cross-view consistency. In *NeurIPS*, 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [54] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015. [2](#), [6](#)
- [55] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. PaMIR: Parametric model-conditioned implicit representation for image-based human reconstruction. *TPAMI*, 2021. [2](#)
- [56] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3D human reconstruction from a single image. In *ICCV*, 2019. [2](#)
- [57] Tiancheng Zhi, Christoph Lassner, Tony Tung, Carsten Stoll, Srinivasa G Narasimhan, and Minh Vo. Texmesh: Reconstructing detailed human texture and geometry from rgb-d video. In *ECCV*, 2020. [2](#)
- [58] Kaiyang Zhou and Tao Xiang. Torchreid: A library for deep learning person re-identification in pytorch. *arXiv preprint arXiv:1910.10093*, 2019. [6](#)
- [59] Luyang Zhu, Konstantinos Rematas, Brian Curless, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Reconstructing nba players. In *ECCV*, 2020. [2](#)