

Generating Texture for 3D Human Avatar from a Single Image using Sampling and Refinement Networks

Sihun Cha¹ , Kwanggyoon Seo¹ , Amirsaman Ashtari¹ , Junyong Noh¹ 

¹Visual Media Lab, KAIST, Republic of Korea




Figure 1: Given a single image, our method generates a texture map by synthesizing textural patterns in the invisible regions of the source image as well as aligning the texture to the surface of the geometry. The top row shows the source human images and the bottom row shows the rendered images of the 3D human avatars with generated texture maps. The 3D human mesh was obtained using Tex2Shape [APTMI9] and the images were sampled from SHHQ[FLJ*22a] dataset.

Abstract

There has been significant progress in generating an animatable 3D human avatar from a single image. However, recovering texture for the 3D human avatar from a single image has been relatively less addressed. Because the generated 3D human avatar reveals the occluded texture of the given image as it moves, it is critical to synthesize the occluded texture pattern that is unseen from the source image. To generate a plausible texture map for 3D human avatars, the occluded texture pattern needs to be synthesized with respect to the visible texture from the given image. Moreover, the generated texture should align with the surface of the target 3D mesh. In this paper, we propose a texture synthesis method for a 3D human avatar that incorporates geometry information. The proposed method consists of two convolutional networks for the sampling and refining process. The sampler network fills in the occluded regions of the source image and aligns the texture with the surface of the target 3D mesh using the geometry information. The sampled texture is further refined and adjusted by the refiner network. To maintain the clear details in the given image, both sampled and refined texture is blended to produce the final texture map. To effectively guide the sampler network to achieve its goal, we designed a curriculum learning scheme that starts from a simple sampling task and gradually progresses to the task where the alignment needs to be considered. We conducted experiments to show that our method outperforms previous methods qualitatively and quantitatively.

CCS Concepts

• **Computing methodologies** → Computer graphics; Computer vision; Image manipulation;

1. Introduction

The demand for animatable 3D human avatars is increasing in various VR/AR applications such as virtual try-on, metaverse, and games. To create an animatable 3D human avatar, it is essential to produce a 3D model that resembles the shape and appearance of the source human appearance. Furthermore, the 3D model should be rigged for animation. These processes often require manual work from artists or rely on a special capture system such as multi-view camera sets or 3D scanners. To alleviate these conditions, numerous methods have been proposed to reconstruct a 3D human avatar [KPD19; APTM19; LCL*19; LIP19; XCM*21; NSH*19; SHN*19; SSSJ20; HCJS20; HXL*20; HXS*21; AZS22] from a single image. In contrast to the reported successes in reconstructing body shapes and poses, restoring the occluded texture for 3D human avatars has been relatively less studied.

Generating a texture map for a 3D human avatar from a single image is challenging due to the following two reasons. First, only portions of the texture information are available from the source image. This is caused by various poses and shapes of the human body and the diverse camera positions. Second, the generated human texture map needs to be semantically aligned with the surface of the target 3D human mesh. As the texture coordinates correspond to the surface geometry of a 3D human mesh, misalignment of the texture can produce a distorted human appearance in a rendered image.

Due to these challenges, generating a 3D human texture map cannot be simply posed as an image inpainting task [LRS*18]. Unlike image painting, where inputs and outputs are spatially aligned, the alignment of inputs and outputs is not guaranteed for 3D human texture generation tasks. Because UV alignment is essential, a process of correcting the spatial structure of the input during the 3D human texture generation task is required. The difference between the image inpainting and 3D human texture generation is highlighted in Figure 2. Therefore, direct application of image inpainting methods to 3D human texture generation tends to result in a misaligned texture map. On the other hand, methods that utilize image-to-image translation [IZZ17] may successfully produce a texture map that is semantically aligned with the UV space of the target mesh [LIP19; WZL*19]. Unfortunately, these CNN-based models tend to learn an average texture from the training data, leading to a blurry result.

In this paper, we propose a method that generates a complete human texture map from a single image while synthesizing the occluded texture with relevance to the given visible appearance. Using a neural network based on sampling and refinement strategies, our method preserves the details given in the source image in the generated texture while retaining the structural alignment with the surface of the target mesh. Similar to previous methods [APT19; LIP19], we convert the source image to a partial texture map and use it as input to our method. We also predict a 3D human mesh based on the SMPL model [LMR*15] from the source image using off-the-shelf method [APT19] and utilize the surface normal information of the 3D human mesh in the sampling process. Given the partial texture map and a normal map, *SamplerNet* completes the missing details by sampling the visible region of the texture and re-arranging them.

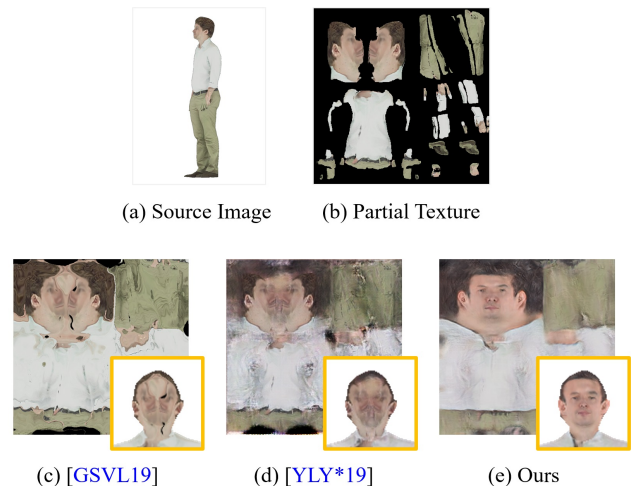


Figure 2: Generated textures using (c) coordinate-based inpainting [GSVL19], (d) color-based inpainting [YLY*19], and (e) our method from (b) partial texture map. The results from inpainting methods show the preservation of the given details and structure of the partial texture map which is created from (a) source image. However, the results failed to align with the surface of the target mesh, which leads to artifacts in rendered images.

The proposed sampler network overcomes the limitation of image inpainting methods [GSVL19; YLY*19], such as the generation of misaligned textures, by learning to align the texture to the corresponding surface of the target 3D human mesh. To guide the *SamplerNet* for effective sampling, we adopt a curriculum learning scheme. Given the sampled texture map produced by *SamplerNet*, *RefinerNet* generates a refined texture and a blending mask. The mask is used for blending the refined texture with the sampled texture to produce the final result. The blending mask helps to preserve the appearance detail presented in the source image and therefore allows to generate the final texture map with improved quality by removing artifacts from the sampled texture. Example results of the generated texture map and rendered images are shown in Figure 1. We conducted a set of experiments to show that our approach outperforms the various baseline methods [IZZ17; WLZ*18; LIP19; XL21; ALY*21] in reproducing the details present in the source image and aligning textures to the surface of the target 3D human mesh.

2. Related Work

2.1. 3D Human Texture Generation with Multi-View Images

There have been many studies that allow to generate 3D human avatars with texture from multi-view images [AMX*18a; AMX*18b; BTTP19; ZLT*20; MAP20; BTTP19]. These methods project the given multi-view images or video frames back to the predicted mesh to create the partial texture maps and combine them with blending techniques to produce the final texture map. However, with only a single view, it is not clear how to extend the methods to capture all the detailed information and fill in the invisible regions.

2.2. 3D Human Texture Generation from a Single Image

Progress in generating 3D human avatars from a single image has shown a big leap in reconstructing body shapes and poses as much as that interest. However, synthesizing a texture that maintains given appearance details and restoring the occluded region with the relevance of the visible region is still challenging. Similar to 3D human texture generation with multi-view images, some methods [NSH*19; SHN*19; HXL*20; HXS*21] acquire the texture by predicting the back view from the given frontal view image followed by projecting the both front and back view images back to the predicted 3D human mesh. While these methods are effective in utilizing the visible textures in the source image, the results are often blurry and poorly reconstructed in the occluded area.

Some studies [GSVL19; LIP19] employ a predefined mapping process that converts the visible human appearance in the source image to the UV space of the SMPL model using DensePose [GNK18]. Lazova et al. [LIP19] utilized the partial texture map produced by a predefined mapping and takes an image-to-image translation approach to generate a full texture map. Because this method learns to generate the missing information based on training data, the generated results can often be blurry and detailed texture patterns are not recovered well when the given garment texture is unseen from the training data.

Instead of using a texture explicitly for supervision, some methods learn the human texture generation in an unsupervised manner, directly producing a texture map from the given image [WZL*19; XL21; CCO22]. These methods can generate a texture map directly from the source image by eliminating the pose variant features and minimizing the identity difference. Xu et al. [XL21] maps the given image to the texture space by predicting flow field and blends it with the generated texture map to eliminate the artifacts and maintain detailed appearance. Inspired from this approach, our method also incorporates blending process with sampled texture and refined texture.

2.3. Pose-guided Image and Video Synthesis

The goal of pose-guided image synthesis is to transfer the person's appearance from a source image to the desired pose. Some approaches [SMX*20; SLGT21; SGLT21] utilize the UV space of the SMPL model as an intermediate representation to achieve the task. These methods generate the latent features, which lie in the SMPL's UV space, and use them to synthesize the image with the given target pose. Other methods [NGK18; LXH*20; YLG*21] directly predict the texture in the UV space to transfer the pose of the source image to the target pose. However, the predicted texture map is used as a reference and undergoes a post-processing step in the image space, which does not consider the alignment with the 3D human mesh. Therefore, the predicted texture may not align with the surface of the target 3D mesh.

Instead of using color pixels from the source image, some methods [GSVL19; ALY*21] map the pixel coordinates of the visible region in the image to the UV space of the SMPL model. The mapped coordinates are then inpainted and used for sampling the source image to create a full texture map. These methods reported better outcomes in retaining the local details presented in

the source image compared to the methods [NGK18] that directly utilize the color pixels of the source image. Similar to these approaches [GSVL19; ALY*21], we use a sampler network to sample the visible regions in the source image to create the texture map. Because the texture map is used for intermediate representation, it is prone to artifacts such as stretching-out or distortion when directly used for rendering as shown in Figure 5. To overcome this and improve quality, we use a texture refinement network that adjusts and refines the produced texture. Detailed illustration and evaluation with previous methods are presented in Section 4.

2.4. Curriculum Learning

After the introduction of curriculum learning by Bengio et al. [BLCW09], the strategy has been adapted to various tasks, such as language modeling [GBM*17], object detection [WWL18], and person re-identification [MMX*17]. We adapt the strategy of curriculum learning by progressively increasing the level of difficulty of the texture sampling task. We train *SamplerNet* from the simplest case of the mapping process and gradually apply geometric augmentation. Through this process, the model learns to complete the texture that is aligned with the surface of the target 3D human mesh while reproducing the detailed appearance presented in the source image.

3. Methods

In this section, we describe the proposed networks and training process. We first preprocess the human appearance in the source image to be mapped into the UV space of a 3D human mesh [LMR*15] and produce a partial texture map and a corresponding visibility mask. From the partial texture map, we obtain the geometry information using Tex2shape [APTM19], which generates the 3D human mesh from a single image by predicting the surface normal and vertex displacements in the UV space of the SMPL model. Given the partial texture map, visibility mask, and normal map, *SamplerNet* samples the missing appearance details from the visible regions of the source image and completes the texture map. In the following process, *RefinerNet* generates a blending mask and a refined texture. Insufficient details and distortion artifacts in the sampled texture produced by *SamplerNet* are adjusted in the refined texture. The final texture is then generated by blending the sampled texture with the refined texture using the blending mask. An overview of the proposed method is shown in Figure 3.

3.1. Preprocessing

Similar to previous methods [LIP19; APTM19; GSVL19; ALY*21], we map the source image into the UV space of the SMPL model. We denote the mapped partial texture as T_{source} . Following Albahar et al. [ALY*21], we combine T_{source} and T_{source}^{mirror} , a semantically mirrored texture map of T_{source} , to produce a symmetric texture map T_{input} . T_{input} is expressed as follows:

$$T_{input} = T_{source} + (T_{source}^{mirror} \odot (1 - M_{source})), \quad (1)$$

where \odot is the Hadamard product. M_{source} is a binary mask in which 1 indicates a valid pixel in T_{source} and 0 indicates an invalid pixel.

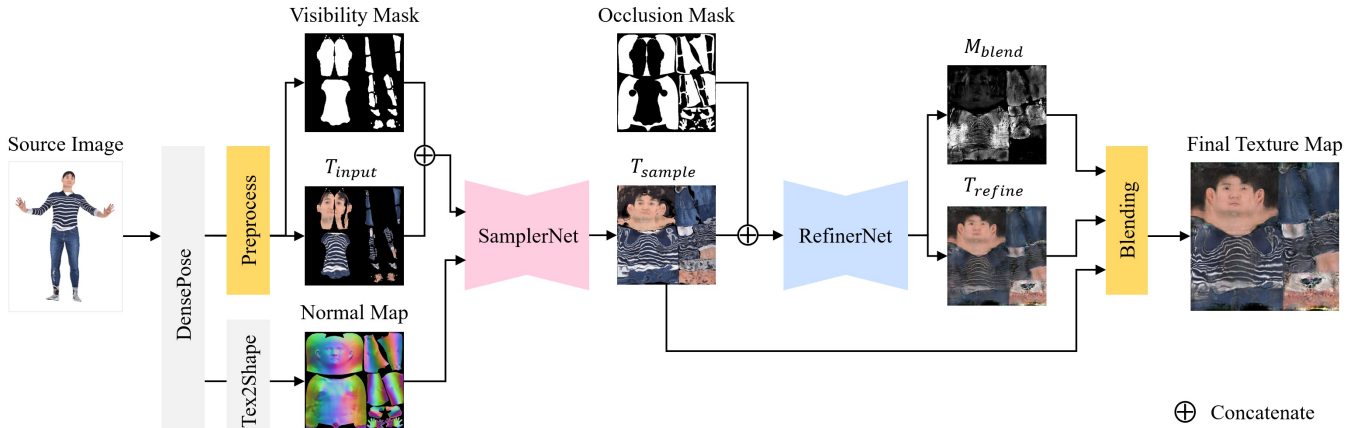


Figure 3: Overview of the proposed method. The source image is processed to create partial texture, visibility mask, and normal map which are given as an input to *SamplerNet*. *SamplerNet* predicts a sampling grid that is used for producing the sampled texture. *RefinerNet* receives the sampled texture and occlusion mask as input, and generates a refined texture and blending mask. The final output is produced by alpha blending the sampled texture with the refined texture using the blending mask.

For the mapping, we use the indexed UV coordinates (IUV) predicted by DensePose [GNK18] to map the source image to the UV space of the SMPL model. The IUV establishes the correspondence between the human appearance in the image and the surface of the SMPL model. The surface is labeled with an index which indicates a predefined body part. Using the IUV and a lookup table provided by DensePose, the visible human appearance in the source image can map to the UV space. Because DensePose divides the surface of the SMPL model to exploit the left-right symmetry, the mirrored texture map T_{source}^{mirror} can be created by switching the index for the IUV of the symmetrical body parts during the mapping process.

3.2. Sampler Network

Given T_{input} , visibility mask, and normal map as input, *SamplerNet* predicts a sampling grid to produce a complete texture map, T_{sample} . The visibility mask is a binary mask that indicates the valid pixel in T_{input} . Using the predicted sampling grid, the visible region of T_{input} is sampled to synthesize the occluded region in T_{input} . Furthermore, the texture in the visible region is re-sampled to be structurally aligned to match the semantic meaning of the UV space. *SamplerNet* consists of two encoders and a single decoder which is similar to the network proposed in Yoon et al. [YLG*21]. The geometry and appearance features are extracted from each dedicated encoder and are fed to the decoder. The appearance features extracted from T_{input} and visibility mask at each layer of the encoder are skip-connected to the corresponding layer in the decoder.

Both encoders used for *SamplerNet* consist of one convolutional layer followed by five layers of residual blocks. The decoder uses the same number of residual block layers with upsampling and one convolutional layer at the end. All activations are LeakyReLU with instance normalization. The residual blocks consist of two convolutional layer with gated convolutions [YLY*19].

When training *SamplerNet*, we observed that *SamplerNet* often

fails to preserve the given appearance detail in T_{input} . To overcome this, we apply the curriculum learning strategy. Before describing the details of the proposed curriculum learning scheme, we will first address the data preparation process for the training.

3.2.1. Training Data Preparation

SamplerNet is trained with a curriculum learning scheme, which begins from the simplest case where the given input texture is perfectly aligned with the UV space, to the hard case where the input texture is misaligned with the UV space. For the simplest case, if the predicted IUV perfectly aligns with the surface coordinates of the SMPL model, partial texture T_{source} can be acquired by masking out the ground truth texture map T_{GT} with M_{source} :

$$T_{source} = T_{GT} \odot M_{source}. \quad (2)$$

We now denote T_{source} obtained by Equation 2 as T_{GT}^M .

Due to imperfect prediction, however, DensePose often fails to locate the exact pixel position that corresponds to the geometric position. This prevents the partial texture from being semantically aligned with T_{GT} . We approximate this misalignment using a geometric transformation function $f(\cdot)$, and thus T_{source} can be expressed as augmentation as follows:

$$T_{source} \approx T_{Augment} = f(T_{GT}^M, \alpha), \quad (3)$$

where α is a control parameter for $f(\cdot)$. By changing α , interpolation from T_{GT}^M to T_{source} is approximated which enables the curriculum learning scheme.

3.2.2. Region-wise Augmentation

For the augmentation $f(\cdot)$, we use the thin-plate-spline (TPS) transformation in a region-wise manner (Figure 4). There are two reasons why naive augmentation techniques are inadequate for texture generation. First, unlike general images, texture maps have a

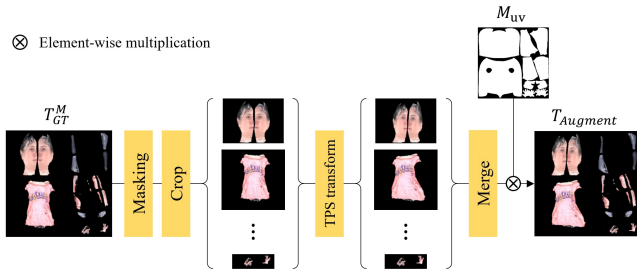


Figure 4: Illustration of the region-wise augmentation.

unique structure that corresponds to a 3D geometry. Second, although T_{source} is non-linearly deformed from T_{GT} , the deformation is limited to the UV structure of the SMPL model. Therefore, augmentation should be performed in consideration of this structure.

We divide the UV space into six different regions which correspond to the head, body, legs, arms, feet, and hands. We crop each region with the bounding box and apply the TPS transform individually. To prevent unintentional cropping, each body part is masked before the process. Then, we merge all of the transformed regions back to form a single texture followed by multiplying it by the UV mask, M_{UV} , to produce the final deformed texture for training. M_{UV} is a binary mask, which represents valid UV coordinates in the UV space. For the transformation, we assign control points to each region and shift these points with a random value determined from a uniform distribution $U(0, \alpha)$. α is expressed as follows:

$$\alpha = \begin{cases} 0, & step = 0, \\ 0.1 + (step * \delta), & step > 0. \end{cases} \quad (4)$$

$step$ indicates the current curriculum step and δ is a hyper-parameter.

3.2.3. Curriculum Learning

The goal of curriculum learning is to guide *SamplerNet* to maintain the given appearance detail in T_{input} while enforcing the semantic alignment. To this end, we divide the training into two curriculum steps: sampling the occluded texture from the visible region and re-arranging the given visible region. First we will describe the data preparation for the curriculum learning.

In the initial step, $step = 0$, the objective is to complete the given partial texture map without considering the alignment in the visible region. Hence, α in Equation 3 is set to 0, which makes $f(\cdot)$ an identity mapping function, resulting in $T_{Augment}$ identical to T_{GT}^M . As T_{GT}^M is semantically aligned with the UV space, *SamplerNet* is encouraged to sample for the missing regions only. A single $step$ is set to 4,000 iterations, which is sufficient for the model to see the whole training data set twice.

After the initial step, α is set to 0.1 and increased by δ after every single step. Here, the goal is to encourage the model to begin re-sampling the visible region to enforce the structural alignment. For the steps equal to or greater than 3, we additionally use the partial texture produced by DensePose to reduce the domain gap between training data and inference data.

The data used in the curriculum learning can be expressed as follows:

$$T_{source} = \begin{cases} T_{Augment}, & step < 3, \\ T_{Augment} \text{ Or } T_{DensePose}, & step \geq 3. \end{cases} \quad (5)$$

where $T_{DensePose}$ is a partial texture produced by DensePose. The variation of the partial texture is visualized in the supplementary material.

The effect of using this augmentation is verified in Section 4.2. With the curriculum learning and the region-wise augmentation, *SamplerNet* effectively samples a partial texture map to synthesize all the occluded region while retaining the given visible part and aligning texture to the surface of the target 3D mesh.

3.2.4. Loss Functions

To enforce the alignment and proper sampling, we train *SamplerNet* by minimizing a reconstruction loss and a perceptual loss between T_{sample} and T_{GT} . The reconstruction loss is expressed as follows:

$$\mathcal{L}_{Recon} = \sum_{i=1}^N ||w_i \cdot M_{body}^i \odot (T_{sample} - T_{GT})||_1, \quad (6)$$

where M_{body}^i is a set of binary masks for the six body parts in the UV space and w_i is its corresponding weights. w_i is set to 6, 1, 1, 1, 1, and 1 for each face, body, leg, arm, foot, and hand.

For the perceptual distance, we use LPIPS [ZIE*18], which can be expressed as follows:

$$\mathcal{L}_{LPIPS} = LPIPS(T_{sample}, T_{GT}), \quad (7)$$

where $LPIPS(\cdot)$ extracts features from the two inputs using AlexNet [Kri14] and calculates the cosine distance between the extracted features.

The total loss function for *SamplerNet* is expressed as follows:

$$\mathcal{L}_{Sampler} = \lambda_{Recon} \mathcal{L}_{Recon} + \lambda_{LPIPS} \mathcal{L}_{LPIPS}, \quad (8)$$

where λ_{Recon} and λ_{LPIPS} are set to 1 and 1, respectively.

3.3. Refiner Network

After *SamplerNet* completes the partial texture, the resulting texture map is refined. *SamplerNet* implicitly learns to sample the missing texture information from the visible region by following the guidance of given geometry information and minimizing the loss in the training process. During the process of sampling the given information according to the geometry information, some artifacts are accompanied as shown in Figure 5.

RefinerNet alleviates this problem by refining the details in T_{sample} . *RefinerNet* employs a U-Net-like architecture with three down and up sampling layers, and 9 residual blocks for the bottleneck. *RefinerNet* receives T_{sample} and occlusion mask as input and produces a refined texture map T_{refine} and a blending mask M_{blend} . The occlusion mask is a binary mask, which is acquired by subtracting M_{src} from M_{UV} .

As described in Section 2.2, Xu et al. [XL21] proposed a mask-fusion strategy that blends the two texture outputs using a predicted




Figure 5: Artifact due to the SamplerNet's prediction error

mask to preserve the fine texture details. Inspired by this, we adapt a texture blending process that exploits details in T_{sample} . The blending mask helps to preserve the given appearance detail by replacing the artifacts in the sampled texture with the refined texture. The final output is computed as follows:

$$T_{final} = T_{sample} \odot M_{blend} + T_{refine} \odot (1 - M_{blend}). \quad (9)$$

Despite the well reconstructed human appearance features, the fine details observed in T_{sample} can sometimes be lost and become blurry in T_{refine} . The blending of T_{refine} and T_{sample} using M_{blend} leverages the advantages of each generated texture. The effect of texture blending is demonstrated in Figure 11.

3.3.1. Loss Functions

To produce a texture map of perceptually plausible quality with minimal artifacts, we minimize the following objective terms when training *RefinerNet*:

$$\begin{aligned} \mathcal{L}_{Refiner} = & \lambda_{Recon} \mathcal{L}_{Recon} + \lambda_{VGG} \mathcal{L}_{VGG} \\ & + \lambda_{GAN} \mathcal{L}_{GAN} + \lambda_{FM} \mathcal{L}_{FM}, \end{aligned} \quad (10)$$

where \mathcal{L}_{Recon} , \mathcal{L}_{VGG} , \mathcal{L}_{GAN} , and \mathcal{L}_{FM} are the reconstruction loss, perceptual loss, adversarial loss, and feature matching loss, respectively.

For \mathcal{L}_{Recon} , instead of calculating the loss between T_{final} and T_{GT} , we calculate the loss between T_{final} and T_{sample} . Because there can be multiple texture maps corresponding to one source image, calculating the loss directly from the ground truth map is restrictive. For example, in the source image, a person wearing a T-shirt with a pattern on the front, may have the same pattern, different pattern, or even no pattern on the back. Thus, a pixel-wise loss with T_{GT} will lead to a texture with the averaged color output (Figure 6). We used the loss between T_{final} and T_{sample} to guide the model in a direction that more respects the estimated T_{sample} . \mathcal{L}_{Recon} is expressed as follows:

$$\mathcal{L}_{Recon} = \|T_{sample} - T_{final}\|_1. \quad (11)$$

For \mathcal{L}_{VGG} , we use the pre-trained VGG-19 [SZ14] to calculate the perceptual distance between T_{final} and T_{GT} by extracting features from each layer l as performed in Wang et al. [WLZ*18]. \mathcal{L}_{VGG} is expressed as follows:

$$\mathcal{L}_{VGG} = \sum_{i=1}^N \frac{1}{w_i} \|VGG_i(T_{GT}) - VGG_i(T_{final})\|_1. \quad (12)$$

Here, VGG_i denotes the layer of the VGG-19 network, where $i \in \{1, 6, 11, 20, 29\}$. w_i is set to 32, 16, 8, 4, and 1 for each layer.



Figure 6: Visual comparison of applying the reconstruction loss to T_{sample} and the ground truth texture map.

We use an adversarial loss with the PatchGAN discriminator [IZZE17]. The objective function can be expressed as follows:

$$\mathcal{L}_{GAN} = \mathbb{E}_{T_{GT}} [\log D(T_{GT})] + \mathbb{E}_{T_{sample}} [\log (1 - D(T_{final}))], \quad (13)$$

where D denotes the discriminator. We additionally use a feature matching loss \mathcal{L}_{FM} [WLZ*18], which minimizes the feature distance between T_{final} and T_{GT} of the discriminator to stabilize the training. \mathcal{L}_{FM} is expressed as follows:

$$\mathcal{L}_{FM} = \sum_{i=1}^N \|D_{l_i}(T_{GT}) - D_{l_i}(T_{final})\|_1, \quad (14)$$

where l_i represents a set of layers after the activation function and $i \in \{1, 2, 3\}$. The weights used in the total loss function are $\lambda_{Recon} = 10$, $\lambda_{VGG} = 10$, $\lambda_{GAN} = 1$, and $\lambda_{FM} = 10$.

3.4. Training Details

We used the Adam optimizer [KB14] with a learning rate of 0.0002 and beta parameters set to 0.9 and 0.999 for both networks, *SamplerNet* and *RefinerNet*. The region-wise augmentation for training *SamplerNet* is applied with the probability of 0.8 and δ is set to 0.025. We used color augmentation for training *RefinerNet* with the probability of 0.5. The batch size is set to 8, and each model is trained separately for 30,000 iterations on an NVIDIA GTX 1080 Ti GPU.

4. Experiments

In this section, we compare our method with previous approaches. Furthermore, several experiments were carried out to evaluate the effect of region-wise augmentation, curriculum learning, and texture blending. A total of 1,441 texture maps were used for training the networks: 929 from SURREAL [VRM*17] dataset, 512 from our newly gathered dataset. The textures were rendered with SMPL [LMR*15] in various human poses to generate partial texture maps for training. To obtain various poses, we randomly sampled 100 poses in the collected animations from Mixamo [Ado]. The textures and poses were randomly paired and rendered in the range of $[-90^\circ, 90^\circ]$ with an interval of 10° . The texture maps were resized to 256×256 for both training and testing. For the evaluation, we used Digital Wardrobe [BTTP19] dataset, which contains 96 textures that are different from the training set. Additionally, we



Figure 7: Visual comparison with previous methods using Digital Wardrobe [BTTP19] dataset. The texture maps are generated for the person viewed from randomly rotated perspectives in the horizontal direction.




Figure 8: Visual comparison with previous methods using Digital Wardrobe [BTTP19] dataset. The images are rendered with texture maps that are generated using the person viewed from the randomly rotated perspectives in the horizontal direction. We used Tex2Shape [APTM19] to estimate the target mesh from the source image.

generated texture using samples from SHHQ [FLJ*22b] dataset to assess the generalizability of our method to real human images. The results are shown in Figure 9.

4.1. Comparisons

We compared our method with previous methods in both generated texture map and rendered image to quantitatively measure the quality of the generated texture map. To evaluate the quality of the generated texture map, we used the following metrics: Structural Similarity (SSIM) [WBSS04], Peak Signal-to-Noise Ratio

(PSNR), and Learned Perceptual Image Patch Similarity (LPIPS) [ZIE*18]. SSIM and PSNR measure the reconstruction quality and LPIPS measures the perceptual quality using VGG-16 [SZ14] as a backbone. We further evaluated the results after applying the textures to the 3D human mesh. For this, we employed LPIPS, and cosine similarity (CosSim) of features extracted from PCB network [SZY*18]. PCB network [SZY*18] is used for a person re-identification task, which aims to find an identical person from the different cameras. The similarity between the source image and the rendered image using the generated texture is high when the CosSim is close to one. To evaluate the quality of the synthesized tex-

Table 1: Quantitative evaluation of the generated texture map. The texture maps were generated using the input images with various views. For the *Avg, we used all the images in the range of $[-90^\circ, 90^\circ]$ with 10° interval as input and averaged the calculated scores from the generated texture maps.

Angle Method	0° (front view)			90° (side view)			*Avg		
	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑
Isola et al. [IZZE17]	0.3142	18.49	0.5305	0.3163	18.48	0.5309	0.3127	18.60	0.5314
Wang et al. [WLZ*18]	0.3085	18.21	0.5323	0.3115	18.40	0.5356	0.3078	18.34	0.5334
Lazova et al. [LIP19]	0.3047	18.91	0.5604	0.3180	18.73	0.5522	0.3075	18.97	0.5561
Xu et al. [XL21]	0.3713	15.70	0.4941	0.3777	15.35	0.4948	0.3755	15.54	0.4929
Albahar et al. [ALY*21] (CCM)	0.2884	16.57	0.5144	0.2700	17.40	0.5336	0.2804	17.15	0.5182
Ours (SamplerNet)	0.2435	17.53	0.5501	0.2445	17.84	0.5620	0.2448	17.78	0.5509
Ours (SamplerNet + RefinerNet)	0.2230	18.04	0.5853	0.2227	18.62	0.5988	0.2236	18.26	0.5837

Table 2: Quantitative evaluation of the rendered image. We compared the methods using the input images with various views and rendered the generated texture maps on the target 3D mesh. We used Tex2Shape [APTM19] to estimate the target mesh from the input image and applied it for rendering in all methods. For the *Avg, we used all the rendered images in the range of $[-90^\circ, 90^\circ]$ with 10° interval as input and averaged the calculated scores.

Angle Method	0° (front view)			90° (side view)			*Avg		
	CosSim-A↑	CosSim-I↑	LPIPS↓	CosSim-A↑	CosSim-I↑	LPIPS↓	CosSim-A↑	CosSim-I↑	LPIPS↓
Isola et al. [IZZE17]	0.7421	0.7685	0.2017	0.7088	0.7405	0.2026	0.7294	0.7577	0.1999
Wang et al. [WLZ*18]	0.7266	0.7542	0.2058	0.6978	0.7279	0.2063	0.7155	0.7430	0.2035
Lazova et al. [LIP19]	0.7510	0.7786	0.1965	0.7199	0.7452	0.2024	0.7384	0.7642	0.1967
Xu et al. [XL21]	0.7480	0.7711	0.1923	0.7338	0.7401	0.1980	0.7421	0.7578	0.1952
Albahar et al. [ALY*21] (CCM)	0.7638	0.7881	0.1944	0.7468	0.7532	0.2007	0.7592	0.7751	0.1952
Ours (SamplerNet)	0.7673	0.7944	0.1931	0.7500	0.7555	0.2007	0.7609	0.7779	0.1954
Ours (SamplerNet + RefinerNet)	0.7678	0.7938	0.1868	0.7504	0.7532	0.1990	0.7615	0.7759	0.1912

ture in the occluded region, CosSim was measured in two different images: one rendered with a pose identical to that of the person in the given image and the other rendered with A-pose in the frontal view. We denote the cosine similarity to the image rendered in the identical pose as CosSim-I and to the image rendered with A-pose as CosSim-A. LPIPS was calculated with the image rendered in the identical pose. To render the image using a generated texture, we obtained a 3D human mesh from the given image using Tex2Shape [APTM19]. We also used RSC-Net [XCM*21] to predict the 3D human pose and the camera parameters from the given image.

We evaluate our method with previous approaches that can generate texture maps using image-to-image translation: Isola et al. [IZZE17], Wang et al. [WLZ*18], and Lazova et al. [LIP19]. Additionally, we also compared our method with approaches that utilize pixel coordinate information: Xu et al. [XL21] and Albahar et al. [ALY*21]. For Albahar et al. [ALY*21], we exploited their coordinate completion model (CCM) as this can be used for texture generation. We used the available implementations provided by the authors except for Lazova et al. [LIP19], in which we tried to replicate the implementation following the descriptions in the paper. All methods were trained with the same data as our methods except for Xu et al. [XL21]. Because our data was not compatible to train the method, we used the pretrained model released with their code. The input for all baseline methods was the symmetric partial texture map which was produced using Equation 1 except for Xu et al. [XL21] which can generate texture map directly from the given

image. To evaluate the robustness of each method against various view angles, we generated texture maps using the images viewed from different perspectives by horizontally rotating the 3D human mesh in the range of $[-90^\circ, 90^\circ]$.

The quantitative results for texture map and rendered image are shown in Table 1 and Table 2, respectively. The methods that generate a texture map using image-to-image translation [IZZE17; WLZ*18; LIP19] achieved high scores in PSNR. Nonetheless, the produced texture lacked the detailed patterns present in the source image and failed to align with the surface of the target 3D mesh as shown in Figure 7(b)-(d) and Figure 8(b)-(d).

The methods that utilize pixel coordinates [XL21; ALY*21] tend to preserve visible texture pattern in the given image better compared to image-to-image translation methods [IZZE17; WLZ*18; LIP19] as shown in Figure 8(e)-(f). Xu et al. [XL21] directly generate a texture map from the given image by predicting the flow field from the source image to the UV space of the SMPL model. This allows the preservation of appearance detail. However, some artifacts are apparent when the side view image is given as shown in Figure 8(e). Albahar et al. [ALY*21] synthesized the texture pattern that is relevant to the given details but often with severe artifacts which defect the alignment with the surface of the target 3D mesh as shown in Figure 8(f). SamplerNet was able to preserve the given details and sample the occluded texture patterns while maintaining the alignment with the surface of the target 3D mesh. This is reflected by higher values of both CosSim-A and CosSim-I reported

in Table 2. Given the sampled texture with improved alignment, our full method produced visually better results compared to those produced by previous methods in terms of maintaining the appearance details present in the source image and synthesizing missing details exploiting the information from the visible region as shown in Figure 8(h).



Figure 9: Rendered images of 3D human avatars produced using real human images. The images are from SHHQ [FLJ*22b] dataset. The 3D human pose and the camera parameters were predicted from the images using RSC-Net [XCM*21].

4.2. Ablation study

4.2.1. Region-wise Augmentation

The goal of region-wise augmentation is to approximate the transformation caused by DensePose [GNK18]. To evaluate the similarity of the proposed augmentation with DensePose, we compared the results from *SamplerNet* trained with several different augmentation alternatives that are applied in a region-wise manner. The curriculum learning was excluded in this experiment to examine the effect of augmentation. α and the probability for the augmentation is fixed to 0.25 and 0.999, respectively. We also compared *SamplerNet* trained with and without region-wise augmentation using TPS. For the case without it, the control points were set at an equispaced 6×6 grid followed by shifting them according to a random value from a uniform distribution $U(0, \alpha)$. We used same α value for both. As shown in Table 3, *SamplerNet* trained with TPS in a region-wise manner resulted in the scores closest to *SamplerNet* trained with DensePose data.

4.2.2. Curriculum Learning

We evaluated the effectiveness of the curriculum learning by comparing our networks trained with and without it. The comparison results are shown in Table 4. The results from the networks trained with curriculum learning achieved better scores for all of the evaluation metrics. Moreover, the network trained with curriculum learning produced qualitatively superior results in terms of reproducing the given appearance as shown in Figure 10.

Table 3: Comparison between various augmentation techniques. The bold number represents the best score and the underlined number represents the score closest to that produced by the model trained with DensePose.

Augmentation	Region-wise	LPIPS↓	PSNR↑	SSIM↑
Rotate	✓	0.2348	17.99	0.5657
Translate	✓	0.2468	18.01	0.5669
TPS	×	0.2323	18.15	0.5758
TPS	✓	<u>0.2270</u>	<u>18.26</u>	<u>0.5817</u>
DensePose	-	0.2142	18.86	0.6071

Table 4: Ablation study for the method trained with and without curriculum learning.

Method	LPIPS↓	PSNR↑	SSIM↑
w/o curriculum learning	0.2238	18.09	0.5920
w/ curriculum learning	0.2220	18.18	0.5935

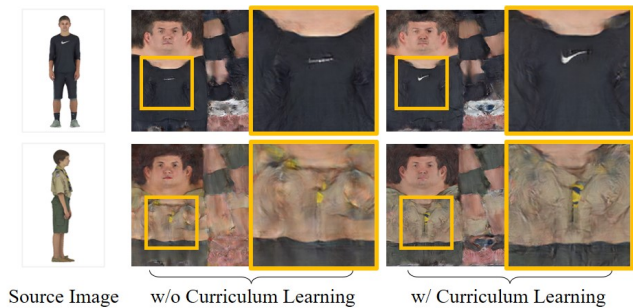


Figure 10: Visual comparison of the results produced with and without curriculum learning.

4.2.3. Texture Blending

We evaluate the effect of blending the refined texture with the sampled texture using a blending mask. We trained *RefinerNet* with two different settings: producing the final texture map using the blending mask and directly generating the final texture map. The results are shown in Table 5 and Figure 11. Although the scores of PSNR and SSIM are similar, our qualitative evaluation (Figure 11) shows improvements in maintaining texture details present in input images, which is reflected in the LPIPS score (Table 5). The model trained without blending fails to produce clear details of the logo or patterns present in the source image as indicated in the orange box.

5. Discussion

5.1. Limitations

Even though our method achieves better results compared to previous methods, we have limitations that need to be addressed. Our method generates textures based on the surface of the target 3D human mesh which is predicted by Tex2Shape [APTM19]. Because Tex2Shape deforms the SMPL [LMR*15] model to obtain the 3D

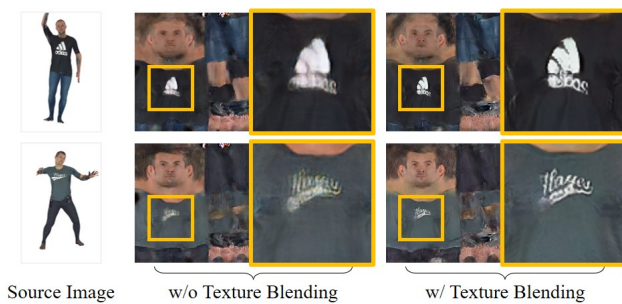


Figure 11: Visual comparison of the resulting texture maps from the model trained with and without using the blending mask.

Table 5: Ablation study for the model trained with and without the texture blending.

Method	LPIPS↓	PSNR↑	SSIM↑
w/o texture blending	0.2276	18.11	0.6007
w/ texture blending	0.2144	18.06	0.5965

human mesh, our model cannot fully handle the textures with loose clothes, such as skirts and long coats.

Similar to other approaches based on deep neural networks, the capability of our method is limited to the training dataset. The size of our training dataset is relatively small compared to previous studies. Thus, our method fails to generate a texture map with various garments such as hats or eye glasses as shown in Figure 7(h) and the last column in Figure 9. Another example is human identity. The SURREAL dataset is limited to a single face-identity, and our newly collected dataset consists of limited ethnicity. This hinders the model from generating a texture map with the personal identity because the same face appeared in the source image. With a bigger and more diverse dataset, this problem will be alleviated.

5.2. Future work

Our method is based on a supervised setting, which requires ground truth data for training. An interesting future research direction is to utilize high-quality image or video datasets, which are relatively easier to acquire, for unsupervised or self-supervised training as exemplified by Grigorev et al. [GII*21]. Another interesting direction is to exploit the high-capability of GAN models [KLA19; FLJ*22b] to generate different views of the source image to achieve better initialization as attempted in video-based methods [AMX*18a; AMX*18b; BTTP19; ZLT*20; MAP20; BTTP19].

6. Conclusion

In this work, we proposed a method for generating a 3D human texture map from a single image. The key idea of our approach is to complete the incomplete partial texture map by using a sampling network followed by adjusting the resulting texture with a refiner network. Compared to previous approaches, our method generates a texture map with improved quality by successfully retaining the

textural patterns presented in the visible regions of the source image while maintaining spatial alignment with the surface of the target 3D human mesh. In addition, our method produces a plausible texture map from a non-frontal view image. We verified the effectiveness of curriculum learning and other design choices we made by ablation studies and showed significant quality improvement in the resulting texture by comparing our approach with previous methods.

Acknowledgements

We thank the anonymous reviewers for their invaluable comments; Gyeong Hun Im for collecting the various pose data; Jaedong Kim for his assistance in implementing the mapping function. This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2020-0-00450, A Deep Learning Based Immersive AR Content Creation Platform for Generating Interactive, Context and Geometry Aware Movement from a Single Image)

References

- [Ado] ADOBE. *Mixamo: Quality 3D character animation in minutes*. <https://www.mixamo.com>. 6.
- [ALY*21] ALBAHAR, BADOUR, LU, JINGWAN, YANG, JIMEL, et al. “Pose with Style: Detail-preserving pose-guided image synthesis with conditional stylegan”. *ACM Transactions on Graphics (TOG)* 40.6 (2021), 1–11 **2, 3, 8**.
- [AMX*18a] ALLDIECK, THIEMO, MAGNOR, MARCUS, XU, WEIPENG, et al. “Detailed human avatars from monocular video”. *2018 International Conference on 3D Vision (3DV)*. IEEE. 2018, 98–109 **2, 10**.
- [AMX*18b] ALLDIECK, THIEMO, MAGNOR, MARCUS, XU, WEIPENG, et al. “Video based reconstruction of 3d people models”. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018, 8387–8397 **2, 10**.
- [APT19] ALLDIECK, THIEMO, PONS-MOLL, GERARD, THEOBALT, CHRISTIAN, and MAGNOR, MARCUS. “Tex2shape: Detailed full human body geometry from a single image”. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019, 2293–2303 **1–3, 7–9**.
- [AZS22] ALLDIECK, THIEMO, ZANFIR, MIHAI, and SMINCHISESCU, CRISTIAN. “Photorealistic Monocular 3D Reconstruction of Humans Wearing Clothing”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022, 1506–1515 **2**.
- [BLCW09] BENGIO, YOSHUA, LOURADOUR, JÉRÔME, COLLOBERT, RONAN, and WESTON, JASON. “Curriculum learning”. *Proceedings of the 26th annual international conference on machine learning*. 2009, 41–48 **3**.
- [BTTP19] BHATNAGAR, BHARAT LAL, TIWARI, GARVITA, THEOBALT, CHRISTIAN, and PONS-MOLL, GERARD. “Multi-garment net: Learning to dress 3d people from images”. *Proceedings of the IEEE/CVF international conference on computer vision*. 2019, 5420–5430 **2, 6, 7, 10**.
- [CCO22] CHANG, SEUNGGYU, CHO, JUNGCHAN, and OH, SONGHWAI. “Texture Generation Using Dual-Domain Feature Flow with Multi-View Hallucinations”. *arXiv preprint arXiv:2203.06901* (2022) **3**.
- [FLJ*22a] FU, JIANGLIN, LI, SHIKAI, JIANG, YUMING, et al. “StyleGAN-Human: A Data-Centric Odyssey of Human Generation”. *arXiv preprint arXiv:2204.11823* (2022) **1**.
- [FLJ*22b] FU, JIANGLIN, LI, SHIKAI, JIANG, YUMING, et al. “StyleGAN-Human: A Data-Centric Odyssey of Human Generation”. *arXiv preprint arXiv:2204.11823* (2022) **7, 9, 10**.

- [GBM*17] GRAVES, ALEX, BELLEMARE, MARC G, MENICK, JACOB, et al. "Automated curriculum learning for neural networks". *international conference on machine learning*. PMLR. 2017, 1311–1320 3.
- [GII*21] GRIGOREV, ARTUR, ISKAKOV, KARIM, IANINA, ANASTASIA, et al. "Stylepeople: A generative model of fullbody human avatars". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, 5151–5160 10.
- [GNK18] GÜLER, RIZA ALP, NEVEROVA, NATALIA, and KOKKINOS, IASONAS. "Densepose: Dense human pose estimation in the wild". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, 7297–7306 3, 4, 9.
- [GSVL19] GRIGOREV, ARTUR, SEVASTOPOLSKY, ARTEM, VAKHITOV, ALEXANDER, and LEMPITSKY, VICTOR. "Coordinate-based texture inpainting for pose-guided human image generation". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, 12135–12144 2, 3.
- [HCJS20] HE, TONG, COLLOMOSSE, JOHN, JIN, HAILIN, and SOATTO, STEFANO. "Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction". *Advances in Neural Information Processing Systems* 33 (2020), 9276–9287 2.
- [HXL*20] HUANG, ZENG, XU, YUANLU, LASSNER, CHRISTOPH, et al. "Arch: Animatable reconstruction of clothed humans". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, 3093–3102 2, 3.
- [HXS*21] HE, TONG, XU, YUANLU, SAITO, SHUNSUKE, et al. "ARCH++: Animation-ready clothed human reconstruction revisited". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, 11046–11056 2, 3.
- [IZZE17] ISOLA, PHILLIP, ZHU, JUN-YAN, ZHOU, TINGHUI, and EFROS, ALEXEI A. "Image-to-image translation with conditional adversarial networks". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, 1125–1134 2, 6, 8.
- [KB14] KINGMA, DIEDERIK P and BA, JIMMY. "Adam: A method for stochastic optimization". *arXiv preprint arXiv:1412.6980* (2014) 6.
- [KLA19] KARRAS, TERO, LAINE, SAMULI, and AILA, TIMO. "A style-based generator architecture for generative adversarial networks". *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2019, 4401–4410 10.
- [KPD19] KOLOTOUROS, NIKOS, PAVLAKOS, GEORGIOS, and DANILIDIS, KOSTAS. "Convolutional mesh regression for single-image human shape reconstruction". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, 4501–4510 2.
- [Kri14] KRIZHEVSKY, ALEX. "One weird trick for parallelizing convolutional neural networks". *arXiv preprint arXiv:1404.5997* (2014) 5.
- [LCL*19] LI, ZHONG, CHEN, LELE, LIU, CELONG, et al. "3D human avatar digitization from a single image". *The 17th International Conference on Virtual-Reality Continuum and its Applications in Industry*. 2019, 1–8 2.
- [LIP19] LAZOVA, VERICA, INSAFUTDINOV, ELДАР, and PONS-MOLL, GERARD. "360-degree textures of people in clothing from a single image". *2019 International Conference on 3D Vision (3DV)*. IEEE. 2019, 643–653 2, 3, 8.
- [LMR*15] LOPER, MATTHEW, MAHMOOD, NAUREEN, ROMERO, JAVIER, et al. "SMPL: A skinned multi-person linear model". *ACM transactions on graphics (TOG)* 34.6 (2015), 1–16 2, 3, 6, 9.
- [LRS*18] LIU, GUILIN, REDA, FITSUM A, SHIH, KEVIN J, et al. "Image inpainting for irregular holes using partial convolutions". *Proceedings of the European conference on computer vision (ECCV)*. 2018, 85–100 2.
- [LXH*20] LIU, LINGJIE, XU, WEIPENG, HABERMANN, MARC, et al. "Neural human video rendering by learning dynamic textures and rendering-to-video translation". *arXiv preprint arXiv:2001.04947* (2020) 3.
- [MAP20] MIR, AYMEN, ALLDIECK, THIEMO, and PONS-MOLL, GERARD. "Learning to transfer texture from clothing images to 3d humans". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, 7023–7034 2, 10.
- [MMX*17] MA, FAN, MENG, DEYU, XIE, QI, et al. "Self-paced co-training". *International Conference on Machine Learning*. PMLR. 2017, 2275–2284 3.
- [NGK18] NEVEROVA, NATALIA, GULER, RIZA ALP, and KOKKINOS, IASONAS. "Dense pose transfer". *Proceedings of the European conference on computer vision (ECCV)*. 2018, 123–138 3.
- [NSH*19] NATSUME, RYOTA, SAITO, SHUNSUKE, HUANG, ZENG, et al. "Siclope: Silhouette-based clothed people". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, 4480–4490 2, 3.
- [SGLT21] SARKAR, KRIPASINDHU, GOLYANIK, VLADISLAV, LIU, LINGJIE, and THEOBALT, CHRISTIAN. "Style and pose control for image synthesis of humans from a single monocular view". *arXiv preprint arXiv:2102.11263* (2021) 3.
- [SHN*19] SAITO, SHUNSUKE, HUANG, ZENG, NATSUME, RYOTA, et al. "Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019, 2304–2314 2, 3.
- [SLGT21] SARKAR, KRIPASINDHU, LIU, LINGJIE, GOLYANIK, VLADISLAV, and THEOBALT, CHRISTIAN. "Humangan: A generative model of human images". *2021 International Conference on 3D Vision (3DV)*. IEEE. 2021, 258–267 3.
- [SMX*20] SARKAR, KRIPASINDHU, MEHTA, DUSHYANT, XU, WEIPENG, et al. "Neural re-rendering of humans from a single image". *European Conference on Computer Vision*. Springer. 2020, 596–613 3.
- [SSSJ20] SAITO, SHUNSUKE, SIMON, TOMAS, SARAGI, JASON, and JOO, HANBYUL. "Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, 84–93 2.
- [SZ14] SIMONYAN, KAREN and ZISSERMAN, ANDREW. "Very deep convolutional networks for large-scale image recognition". *arXiv preprint arXiv:1409.1556* (2014) 6, 7.
- [SZY*18] SUN, YIFAN, ZHENG, LIANG, YANG, YI, et al. "Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline)". *Proceedings of the European conference on computer vision (ECCV)*. 2018, 480–496 7.
- [VRM*17] VAROL, GUL, ROMERO, JAVIER, MARTIN, XAVIER, et al. "Learning from synthetic humans". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, 109–117 6.
- [WBSS04] WANG, ZHOU, BOVIK, ALAN C, SHEIKH, HAMID R, and SIMONCELLI, EERO P. "Image quality assessment: from error visibility to structural similarity". *IEEE transactions on image processing* 13.4 (2004), 600–612 7.
- [WLZ*18] WANG, TING-CHUN, LIU, MING-YU, ZHU, JUN-YAN, et al. "High-resolution image synthesis and semantic manipulation with conditional gans". *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, 8798–8807 2, 6, 8.
- [WWL18] WANG, JIASI, WANG, XINGGANG, and LIU, WENYU. "Weakly-and semi-supervised faster r-cnn with curriculum learning". *2018 24th International Conference on Pattern Recognition (ICPR)*. IEEE. 2018, 2416–2421 3.
- [WZL*19] WANG, JIAN, ZHONG, YUNSHAN, LI, YACHUN, et al. "Re-identification supervised texture generation". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, 11846–11856 2, 3.
- [XCM*21] XU, XIANGYU, CHEN, HAO, MORENO-NOGUER, FRANCESC, et al. "3D human pose, shape and texture from low-resolution images and videos". *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021) 2, 8, 9.

- [XL21] XU, XIANGYU and LOY, CHEN CHANGE. “3D human texture estimation from a single image with transformers”. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, 13849–13858 [2](#), [3](#), [5](#), [8](#).
- [YLG*21] YOON, JAE SHIN, LIU, LINGJIE, GOLYANIK, VLADISLAV, et al. “Pose-guided human animation from a single image in the wild”. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, 15039–15048 [3](#), [4](#).
- [YLY*19] YU, JIAHUI, LIN, ZHE, YANG, JIMEI, et al. “Free-form image inpainting with gated convolution”. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019, 4471–4480 [2](#), [4](#).
- [ZIE*18] ZHANG, RICHARD, ISOLA, PHILLIP, EFROS, ALEXEI A, et al. “The unreasonable effectiveness of deep features as a perceptual metric”. *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, 586–595 [5](#), [7](#).
- [ZLT*20] ZHI, TIANCHENG, LASSNER, CHRISTOPH, TUNG, TONY, et al. “Texmesh: Reconstructing detailed human texture and geometry from rgb-d video”. *European Conference on Computer Vision*. Springer. 2020, 492–509 [2](#), [10](#).