

PARTE: Part-Guided Texturing for 3D Human Reconstruction from a Single Image

Hyeongjin Nam^{*1} Donghwan Kim^{*1} Gyeongsik Moon^{†2} Kyoung Mu Lee¹

¹Dept. of ECE&ASRI, Seoul National University ²Dept. of CSE, Korea University

{namhjsnu28, dh971106, kyoungmu}@snu.ac.kr, mks0601@korea.ac.kr

<https://hyenie1228.github.io/PARTE/>

Abstract

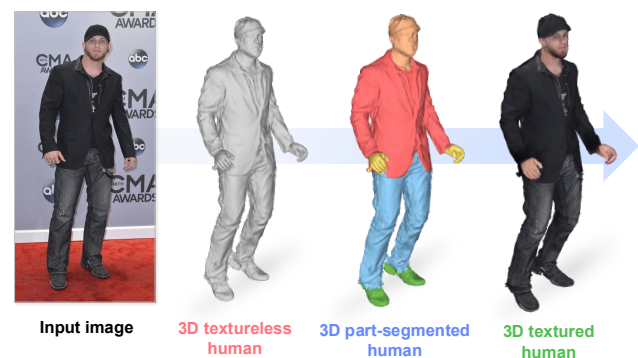
The misaligned human texture across different human parts is one of the main limitations of existing 3D human reconstruction methods. Each human part, such as a jacket or pants, should maintain a distinct texture without blending into others. The structural coherence of human parts serves as a crucial cue to infer human textures in the invisible regions of a single image. However, most existing 3D human reconstruction methods do not explicitly exploit such part segmentation priors, leading to misaligned textures in their reconstructions. In this regard, we present **PARTE**, which utilizes 3D human part information as a key guide to reconstruct 3D human textures. Our framework comprises two core components. First, to infer 3D human part information from a single image, we propose a 3D part segmentation module (PartSegmenter) that initially reconstructs a textureless human surface and predicts human part labels based on the textureless surface. Second, to incorporate part information into texture reconstruction, we introduce a part-guided texturing module (PartTexturer), which acquires prior knowledge from a pre-trained image generation network on texture alignment of human parts. Extensive experiments demonstrate that our framework achieves state-of-the-art quality in 3D human reconstruction.

1. Introduction

3D human reconstruction aims to recover both the geometry and texture of a human from a single image. This task is essential for various applications, such as filmmaking, game development, and AR/VR. The key challenge in 3D human reconstruction is inferring human textures for invisible regions by extracting appearance cues, such as cloth coverage and hair color, from the limited front-view image.

^{*}Equal contribution. [†]Corresponding author.

[‡]This work was done while Hyeongjin Nam was in KRAFTON.



(a) Overview of PARTE



(b) Importance of part alignment

Figure 1. **Overview of PARTE.** Our framework reconstructs a realistic 3D textured human from a single image by effectively improving part alignment with 3D human part information.

Despite remarkable advances in 3D reconstruction of human geometry [20, 48, 62, 73, 74], reconstructing high-fidelity textures has been relatively unexplored. Most existing methods [6, 10, 22, 24, 31, 61, 77, 84, 88] train their networks on 3D human scan datasets [8, 23, 47, 58, 79] to directly predict the colors corresponding to 3D points on the human surface. More recently, several works [2, 30, 37, 80, 82, 85] have leveraged a pre-trained diffusion model [60] that contains rich image prior knowledge to reconstruct realistic human textures. However, a critical challenge remains

in human texture reconstruction: part misalignment. Existing methods often fail to reconstruct textures that accurately correspond to each human part. They primarily rely on the global context of the image rather than explicitly identifying local semantics, which makes it difficult to capture part-specific information from the input image. Furthermore, existing methods often ignore the structural information of the human, resulting in incorrect textures that do not match the underlying structure. As shown in Fig. 1 (b), these failure cases highlight the requirement for part alignment to enhance the overall texture reconstruction quality.

To address the part misalignment, we propose **PARTE**, **PARt**-guided **TE**xting for 3D human reconstruction. This framework estimates 3D part segmentations of the human surface (*e.g.*, upper-cloth and footwear) and utilizes them as main guidance for reconstructing human textures, as shown in Fig. 2. The part segmentations effectively guide subsequent texturing by specifying distinct human part regions, resulting in well-aligned human textures corresponding to the human parts. Our framework consists of two components: 1) **PartSegmenter** and 2) **PartTexturer**. **PartSegmenter** is a module that predicts 3D human part segmentations from an input image. The key challenge of this module lies in predicting part labels for regions not visible in the front-view. To address this, **PartSegmenter** initially reconstructs the textureless 3D human surface and infers the part segmentation based on the textureless surface. The textureless human surface contains structural cues, such as depth variations and human part boundaries, which cannot be captured in the front-view, aiding inference of part labels for invisible regions. By incorporating such geometric cues with the visual feature of the image, **PartSegmenter** effectively resolves ambiguity in invisible regions, enabling accurate and consistent 3D part segmentation.

PartTexturer is a module that reconstructs human textures based on our estimated 3D part segmentations as key guidance. The key challenge of this module is to infer a realistic human texture that satisfies two constraints: alignment with the segmented human parts and visual consistency with the input image. To this end, we devise a specialized diffusion network, which incorporates both the part segmentations and the visual features extracted from the input image. From the part segmentations, the model receives structural guidance, ensuring that textures are assigned accurately to distinct part regions with clear boundaries. From the input image, it captures detailed visual features of each human part, providing exact texture guidance for each respective part. By leveraging the diffusion network’s effectiveness in part-aware texturing, our framework accurately reconstructs human textures corresponding to the part segmentations, thus improving the overall quality and consistency of the reconstructed human textures.

As a result, our proposed PARTE significantly improves

the visual fidelity of reconstructed 3D humans, addressing part misalignment issues, compared to existing state-of-the-art 3D human reconstruction methods. Our contributions can be summarized as follows.

- We propose **PARTE**, a 3D human reconstruction framework that recovers realistic human textures with precise alignment to each human part.
- To acquire precise 3D part segmentation, we present a part segmentation module, **PartSegmenter**, which integrates information from both visible and invisible regions of the input image.
- To accurately reconstruct human textures that align with human parts, we propose a part-guided texturing module, **PartTexturer**, which effectively combines part segmentations and visual features from the input image.

2. Related Works

3D human reconstruction. 3D human reconstruction methods [3, 38, 71] can be categorized into regression- and optimization-based approaches. The regression-based methods [5, 13, 22, 24, 31, 42, 49–51, 53, 61, 62, 67, 88] directly estimate geometry (*i.e.*, distance fields) and texture (*i.e.*, RGB color) from an input image, by training neural networks on 3D human scan datasets. PIFu [61] is proposed to spatially align image features with the input image to learn geometry and texture effectively. SiTH [24] adopts a two-stage approach, first predicting a back-view image, then reconstructing a 3D human using both the front- and back-view images. However, such approaches are fundamentally limited by the scarcity of 3D training data, as acquiring high-quality 3D human scans requires specialized studios or a process of capturing multi-view images. To address this limitation, more recent methods [18, 24, 30, 63, 68, 80, 85] have adopted an optimization-based approach that takes advantage of the prior knowledge of pre-trained diffusion models [60] on large-scale 2D image datasets [16, 44, 45]. Since large-scale 2D images are relatively easy to collect, the optimization-based approach has demonstrated realistic 3D reconstructions [11, 32, 54, 70]. TeCH [30] utilizes a text-to-image diffusion model for reconstruction, ensuring the appearance of the reconstructed 3D human matches the text information of the image. SIFU [85] adopts a coarse-to-fine manner, first reconstructing human texture through a regression-based method and refining it via an optimization-based method.

Recently, several works [20, 75, 82] have utilized human part priors in their 3D human reconstruction pipelines. Textformer [75] and 2K2K [20] use 2D part segmentations from the input image without explicitly considering the overall structure of 3D human parts. HumanRef [82] uses hand-crafted 3D bounding boxes of human parts in the reconstruction, but the bounding boxes do not reflect the actual human parts in the image. Unlike these methods, our pro-

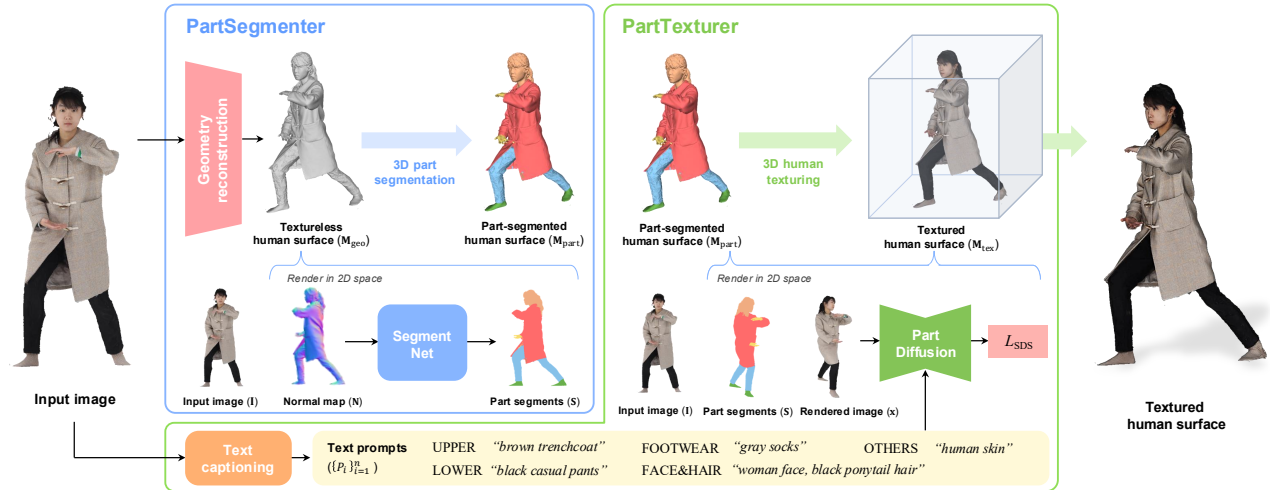


Figure 2. **Overall pipeline of PARTE.** Our framework initially reconstructs a textureless human surface from a single image and textures it based on two core modules. **PartSegmenter** predicts 3D human parts of the textureless human surface by incorporating information from the input image and normal maps of the textureless human surface. Based on the 3D human part segmentations, **PartTexturer** reconstructs the human textures using PartDiffusion network that infers plausible human appearance corresponding to the human parts.

posed PARTE explicitly infers dense 3D part labels in 3D space, providing two key advantages. First, since our framework estimates part segmentations in 3D space, the 3D part segmentations can provide accurate part guidance from any viewpoint instead of relying solely on front-view information. Second, our estimated 3D part segmentations reflect the actual appearance of the input image, ensuring a correct texture assignment to each human part. With these advantages, our PARTE effectively reconstructs human textures with precise alignment across human parts.

3D human generation. Alongside image-based 3D human reconstruction methods, 3D human generation [1, 4, 9, 12, 15, 17, 19, 25–29, 33, 35–37, 43, 48, 72, 76] using different modalities (e.g., text prompt) has also been actively studied. AG3D [15] trains a 3D human generation network with adversarial loss to learn realistic human appearances. Avatar-CLIP [25] proposes to generate 3D humans whose rendered images are conditioned to align with text descriptions using a pre-trained vision-language model [57]. Unlike 3D human generation, image-based 3D human reconstruction has the additional challenge of preserving the appearance of the input image, such as its human part structure. To address this challenge, we propose part-guided texturing that effectively reconstructs human textures while preserving human appearance in the image.

Region-aware image generation. Region-aware image generation [7, 14, 21, 40, 41, 46, 52, 64, 65, 69, 78, 81, 87, 89] aims to synthesize images where each region is generated in alignment with its intended semantics, enabling more structured and controllable image generation. LayoutDiffusion [87] modulates the diffusion network to generate images that conform to text prompts and 2D bounding


boxes. InstanceDiffusion [69] improves region-aware control by conditioning generation on region-specific prompts within corresponding regions. In 3D human reconstruction, the human body naturally consists of distinct regions, human parts, each with its own structural and visual significance. Building on the strong generative power of region-aware image generation, we introduce part-guided texturing, which infers human textures explicitly guided by human part information.

3. PartSegmenter

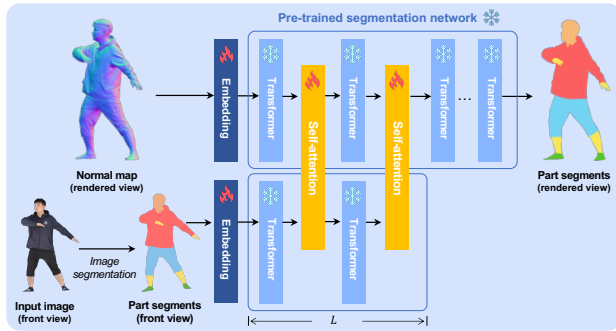
PartSegmenter estimates the 3D part-segmented human surface M_{part} from an input image, where each vertex on the surface is assigned a semantic label of $n = 5$ human parts: face & hair, upper-clothes, lower-clothes, footwear, and others. The acquired 3D part-segmented surface serves as input for the next module, PartTexturer (Sec. 4), which reconstructs human textures.

3.1. Geometry reconstruction

Given an input image, we first reconstruct the high-fidelity geometry M_{geo} of a 3D human surface before predicting the 3D part-segmented human surface M_{part} . For our geometry reconstruction, we employ an off-the-shelf reconstruction method, TeCH [30] that has shown superior performance in real-world scenarios. Note that our framework is compatible with various 3D human geometry reconstruction methods. A detailed explanation of the geometry reconstruction process is provided in the supplementary material.



(a) Overview of PartSegmenter



(b) SegmentNet architecture

Figure 3. **Detailed process of PartSegmenter.** (a) From a textureless human surface, PartSegmenter performs normal rendering and acquires part segments via SegmentNet. The part segments are used to vote for 3D human part labels on the human surface. (b) SegmentNet incorporates features from the normal map and the front-view image for accurate part segmentation.

3.2. 3D part segmentation

Fig. 3 (a) illustrates the process of estimating the 3D segmented human surface \mathbf{M}_{part} from the initially reconstructed textureless human surface \mathbf{M}_{geo} . From the textureless human surface, we render a normal map $\mathbf{N}^{\mathbf{k}}$, where \mathbf{k} is a camera viewpoint sampled uniformly from a spherical distribution. With the normal map, SegmentNet predicts its 2D part segments $\mathbf{S}^{\mathbf{k}}$, assigning one of n part labels to each region. We collect 2D part segments from 30 uniformly distributed viewpoints and aggregate them via voting to construct the part-segmented human surface \mathbf{M}_{part} .

This strategy of aggregating 2D segmentation in 3D space, rather than directly predicting part labels in 3D, enhances generalization. A naive approach to 3D part segmentation is to predict part segments directly in 3D space [55, 56, 66, 86] by training a prediction network on 3D human datasets. However, these 3D datasets require restricted environments, such as capture studios, which limit their diversity and fail to capture real-world scenarios. Consequently, designing a pipeline to directly predict 3D part labels results in poor generalization. In contrast, 2D human datasets [16, 44, 45] are relatively abundant, enabling more

generalizable 2D human part segmentation. To leverage this advantage, we adopt a strategy that first performs segmentation in a 2D space and then aggregates the results in 3D space, achieving accurate 3D part segmentation.

3.3. SegmentNet

Fig. 3 (b) illustrates the detailed architecture of SegmentNet. SegmentNet comprises two processing branches that interact through self-attention mechanisms: one for the normal map (rendered view) and one for the input image (front view). The first branch primarily captures geometric information, such as cloth boundaries, from the normal map. The second branch extracts semantic cues, such as cloth length and style, from the part segments obtained by applying the off-the-shelf image segmentation method (*i.e.*, Sapiens [34]) to the input image. These branches are processed through $L = 10$ Transformer layers, where self-attention layers are applied after each Transformer to aggregate features between the two branches. Then, the first branch originating from the normal map is processed through the remaining Transformer layers to predict the part segments. All Transformer weights are initialized from the pre-trained Sapiens network [34] and are kept frozen during the training of SegmentNet. The training details are provided in the supplementary material.

With this design, SegmentNet effectively incorporates features of two branches, resulting in accurate segmentation results. The normal map (first branch) provides structural details (*e.g.*, human part boundaries) on the rendering viewpoint that include invisible regions from the input image. However, the normal map can be blurry or noisy due to geometric errors in the textureless human surface. Thus, relying solely on the normal map can provide wrong segmentations. To improve the robustness of segmentation, SegmentNet integrates features from visible regions, extracted from the front-view part segments (second branch), which provide semantic cues about human appearance (*e.g.*, cloth style). By incorporating the semantic information from the visible region with geometric information from the invisible region, SegmentNet ensures consistent part segmentation even with an imperfect normal map. As a result, SegmentNet predicts 3D-consistent and reliable human part labels, enabling PartSegmenter to produce an accurate 3D part-segmented human surface.

4. PartTexturer

PartTexturer reconstructs human textures on the textureless human surface \mathbf{M}_{geo} , producing the fully textured 3D human surface \mathbf{M}_{tex} . Based on the part-segmented human surface \mathbf{M}_{part} , we distinctly guide each human part’s texture using our proposed diffusion model, PartDiffusion.

4.1. Text captioning

Before texturing, our framework acquires text prompts that provide descriptive information about the input image. Each prompt describes attributes, such as hairstyle and cloth color, of $n = 5$ human parts. These text prompts serve as guidance signals for the semantic information of human appearance during the reconstruction of part-specific textures. In our framework, we use an off-the-shelf text captioning method, BLIP [39].

4.2. 3D human texturing

To reconstruct the human texture on the textureless human surface \mathbf{M}_{geo} , we optimize an MLP network ϕ that predicts RGB colors from 3D points on the human surface. The overall loss function of the optimization is as follows:

$$\mathcal{L} = \mathcal{L}_{\text{recon}} + \mathcal{L}_{\text{SDS}}, \quad (1)$$


where $\mathcal{L}_{\text{recon}}$ is the reconstruction loss, which is computed as the L2 distance between the input image and the rendered image of the human surface from the front-view. \mathcal{L}_{SDS} is a score distillation sampling (SDS) loss function [54], which leverages a pre-trained diffusion network to infer invisible human textures from non-frontal views. Specifically, at a randomly selected camera viewpoint \mathbf{k} , we render an image $\mathbf{x}^{\mathbf{k}}$ from human surface colors predicted by the MLP network ϕ . Based on the rendered image $\mathbf{x}^{\mathbf{k}}$, the SDS loss is calculated as

$$\nabla_{\phi} \mathcal{L}_{\text{SDS}} = \mathbb{E}[w_t(\hat{\epsilon}_t(\mathbf{x}_t^{\mathbf{k}}; \mathbf{S}^{\mathbf{k}}, \mathbf{I}, \{P_i\}_{i=1}^n) - \epsilon_t) \frac{\partial \mathbf{x}_t^{\mathbf{k}}}{\partial \phi}], \quad (2)$$

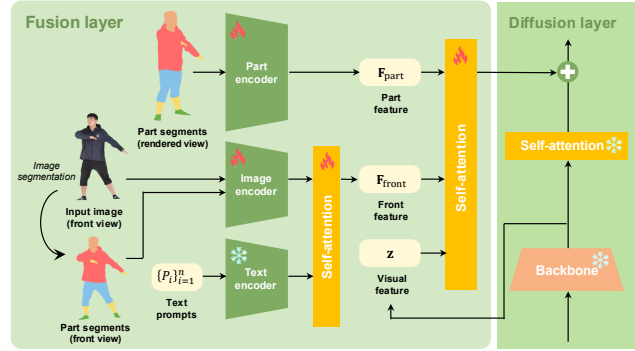
where t denotes the noise level, $\mathbf{x}_t^{\mathbf{k}}$ represents the rendered human surface perturbed by the noise ϵ_t , and w_t is a weighting factor dependent on the noise level t . $\mathbf{S}^{\mathbf{k}}$ denotes the part segments obtained by rendering the part-segmented human surface \mathbf{M}_{part} , \mathbf{I} is the input image, and $\{P_i\}_{i=1}^n$ are text prompts obtained from text captioning. This loss function minimizes the distance between the predicted noise $\hat{\epsilon}_t(\cdot)$ and its counterpart noise ϵ_t . Minimizing the distance enforces the rendered human surface $\mathbf{x}^{\mathbf{k}}$ to match the distribution of human appearances learned by the diffusion network. In this loss function, a key requirement is for the diffusion network to predict accurate noise while maintaining alignment with the conditioning inputs: $\mathbf{S}^{\mathbf{k}}$, \mathbf{I} , and $\{P_i\}_{i=1}^n$. To this end, we propose PartDiffusion, whose detailed explanation is provided in the following section (Sec. 4.3).

4.3. PartDiffusion

Fig. 4 (a) illustrates an overview of PartDiffusion. PartDiffusion takes three conditioning inputs, part segments $\mathbf{S}^{\mathbf{k}}$, input image \mathbf{I} , and text prompts $\{P_i\}_{i=1}^n$ and predicts the noise corresponding to the noisy image $\mathbf{x}_t^{\mathbf{k}}$, where \mathbf{k} is a rendering viewpoint. To integrate these conditions into the



(a) Overview of PartDiffusion



(b) PartDiffusion architecture

Figure 4. **Detailed process of PartTexturer.** (a) To infer texture for a specific rendering view, PartDiffusion integrates three key sources of information: the front-view image, text prompts, and part segments of the rendering view. (b) For effective integration, we adopt self-attention on the separately encoded features.

visual features of a pre-trained diffusion network, we aggregate their features via a fusion layer. All weights of the diffusion network are initialized from StableDiffusion [59] and remain frozen while training PartDiffusion. The training details are provided in the supplementary material.

Fig. 4 (b) shows the detailed process of the fusion layer, which aggregates the three conditioning inputs: part segments (rendered view), input image (front view), and text prompts. The part segments are encoded into a part feature \mathbf{F}_{part} that captures the human part structural information at the rendered viewpoint. The input image is processed through the off-the-shelf image segmentation method, Sapiens [34], to extract its part segments (front-view). Each part region in the part segments is element-wise multiplied with the input image, creating $n = 5$ image patches, each containing the appearance of its corresponding part. The image patches and text prompts are encoded and incorporated separately through a self-attention layer to produce a front feature $\mathbf{F}_{\text{front}}$. This front feature captures the semantic information of each human part in the input image, such as the color and style of the cloth. The extracted part feature \mathbf{F}_{part} and front feature $\mathbf{F}_{\text{front}}$ are aggregated with a visual feature \mathbf{z} from the diffusion layer, through a self-attention

layer. The aggregated features are then added to the output of the diffusion layer. With this architectural design, PartDiffusion effectively integrates the structural information from the part segments in the rendered view and the semantic information from the front-view. As a result, our proposed PartDiffusion accurately infers human textures that are structurally aligned with human parts and visually consistent with the input image, thereby enabling **PartTexturer** to produce high-quality textured human surface \mathbf{M}_{tex} .

5. Experiments

5.1. Datasets

THuman2.1. THuman2.1 [79] consists of high-quality 3D scans of 2,445 humans. To obtain ground truth (GT) part segmentations for the scans, we apply Sapiens [34] to 360 rendered images from uniformly distributed viewpoints and aggregate the results in 3D space. For our experiments, we uniformly sample 80 scans for evaluation, while the remaining scans are used as training data to train SegmentNet in PartSegmenter and PartDiffusion in PartTexturer.

HuMMan. HuMMan [8] is a large-scale dataset containing diverse human poses with 3D scans. To obtain GT part segmentations, we apply Sapiens [34] to 360 rendered images from uniformly distributed viewpoints and aggregate the results in 3D space. For our experiments, we uniformly sample 128 scans from the official test set. This dataset is used only for evaluation purposes.

SHHQ. SHHQ [16] contains a large number of in-the-wild images showcasing diverse cloth styles and human appearances. Since SHHQ does not have GT 3D human data, it is used only for qualitative evaluation.

5.2. Evaluation Metrics.

3D geometry reconstruction. To evaluate the geometry of reconstructed 3D human surfaces, we measure P2S (Point-to-Surface) and CD (Chamfer distance), between the reconstructed and GT human surfaces. These metrics are measured in centimeters.

3D texture reconstruction. To evaluate the texture of reconstructed 3D human surfaces, we employ PSNR (Peak Signal-to-Noise Ratio), LPIPS [83], and Part IoU as evaluation metrics. To compute these metrics, we first render both the reconstructed and GT human surfaces from pre-defined viewpoints at 0° , 90° , 180° , and 270° , where 0° represents the front-view. We then compute PSNR, LPIPS, and Part IoU, between the rendered images. PSNR and LPIPS quantify the visual fidelity and perceptual similarity of the reconstructed human textures. Part IoU evaluates the part alignment of the reconstructed human textures in 2D space. It is computed as the average IoU of part segments of the rendered images, where each segment is obtained by applying Sapiens [34] to the rendered images.


3D part segmentation. To evaluate 3D part segmentation, we introduce two metrics: Part CD and Label Acc_{GT} . Part CD evaluates the part alignment of the reconstructed textures in 3D space. It is computed as the average Chamfer distance across all 3D human parts between the predicted part-segmented human surfaces and their GT counterparts. We measure Part CD in centimeters. Label Acc_{GT} evaluates the accuracy of part segmentation on GT textureless human surface, by eliminating the effect of geometric error. To compute this, we run PartSegmenter on GT textureless human surfaces and compute the proportion of correctly classified part labels on the surface. Note that Label Acc cannot be calculated based on reconstructed surfaces, as the reconstructed surfaces have different topologies from the GTs.

2D image generation. To verify the learned image priors of the diffusion network used in PartTexturer, we evaluate generated images by the diffusion network, using PSNR and Part IoU. For evaluation, we obtain 2D part segments by rendering GT human surfaces at pre-defined viewpoints 0° , 90° , 180° , and 270° , where 0° represents the front-view. Subsequently, we generate 10 human images corresponding to each 2D part segment. We compute PSNR and Part IoU between the generated images and their GT counterparts, which are the rendered images from the GT human surfaces. PSNR evaluates whether the generated images preserve the front-view image’s appearance across various viewpoints. Part IoU evaluates whether the generated images align with the conditioned 2D part segments.

5.3. Ablation study

We carry out all ablation studies on THuman2.1 [79] evaluation set. The geometry reconstruction results are the same across all ablation studies, with P2S and CD being 2.984 and 3.008, respectively.

Effectiveness of PartSegmenter. Fig. 5 (a) and Tab. 1 demonstrate the effectiveness of PartSegmenter in both 3D part segmentation and texture reconstruction. As shown in the first block of Tab. 1, part segmentation information is highly beneficial for texture reconstruction. Without part segmentation, where the entire human silhouette is treated as a single part segment, the PartSegmenter fails to distinguish different human parts, leading to texture misalignment. On the other hand, our framework leverages part segmentation as core guidance for texture reconstruction, resulting in accurate textures that are correctly aligned with their corresponding human parts. Furthermore, our PartSegmenter achieves more accurate 3D part segmentation compared to the existing 3D human segmentation method, Human3D [66]. For a fair comparison, we fine-tuned Human3D on the same training dataset used for PartSegmenter. Human3D takes a human point cloud as input without utilizing image information. However, using only geometric information results in imprecise segmentation, as it intro-



(a) Effectiveness of PartSegmenter



(b) Effectiveness of PartDiffusion

Figure 5. **Impact of (a) PartSegmenter and (b) PartDiffusion in 3D human texture reconstruction.**

duces ambiguity in capturing semantic cues, such as cloth style. On the other hand, our PartSegmenter utilizes front-view image information along with geometric information (*i.e.*, normal map), effectively capturing semantic cues from the image to support accurate segmentation. Thus, PartSegmenter achieves superior performance in part segmentation, which contributes to accurate human texture reconstruction. The second block of Tab. 1 shows that incorporating part segments from the front-view image leads to more accurate part segmentation by SegmentNet. The normal map, which serves as the input to SegmentNet, is often noisy. To compensate for missing or inaccurate information in the noisy input, our SegmentNet utilizes part segments from the front-view image, which provide additional structural cues. Accordingly, SegmentNet achieves high part segmentation accuracy and leads to accurate texture reconstruction.

Effectiveness of PartTexturer. Fig. 5 (b) and Tab. 2 demonstrate the effectiveness of PartTexturer in texture reconstruction. As shown in the first block of Tab. 2, our proposed PartDiffusion in PartTexturer is much more effective in both 2D image generation and 3D human texturing than other diffusion networks. StableDiffusion [59] does not use part segmentation and relies solely on text prompts as its conditioning input. Without the guidance of part segmentation, it fails to infer the correct textures corresponding to human parts, leading to misaligned texture reconstruction. InstanceDiffusion [69], while taking human segments as input, does not incorporate the information from the front-view image. Accordingly, it often reconstructs human textures inconsistent with the front-view image’s appearance.

Methods	3D part segmentation		3D texture reconstruction		
	Part CD \downarrow	Label Acc \uparrow_{GT}	PSNR \uparrow	LPIPS \downarrow	Part IoU \uparrow
* Effectiveness of PartSegmenter					
w/o part segmentation	-	-	20.992	0.115	0.461
Human3D \dagger [66]	5.231	0.881	21.043	0.113	0.574
PartSegmenter (Ours)	4.369	0.947	22.175	0.096	0.641
* Variations of SegmentNet design					
w/o front-view image	5.087	0.941	21.654	0.105	0.635
SegmentNet (Ours)	4.369	0.947	22.175	0.096	0.641

Table 1. **Ablation studies for PartSegmenter.** \dagger indicates that the method is fine-tuned on our training datasets for fair comparison.

Methods	2D image generation		3D texture reconstruction		
	PSNR \uparrow	Part IoU \uparrow	PSNR \uparrow	LPIPS \downarrow	Part IoU \uparrow
* Effectiveness of PartDiffusion					
StableDiffusion [59]	10.719	0.138	19.276	0.166	0.472
InstanceDiffusion [69]	16.125	0.563	20.860	0.115	0.527
PartDiffusion (Ours)	20.109	0.854	22.175	0.096	0.641
* Variations of PartDiffusion design					
w/o front-view image	19.203	0.845	21.005	0.110	0.580
w/o text prompts	19.422	0.803	20.859	0.114	0.588
PartDiffusion (Ours)	20.109	0.854	22.175	0.096	0.641

Table 2. **Ablation studies for PartTexturer.**

On the other hand, our proposed PartDiffusion explicitly conditions the texturing process with both part segmentation and front-view image, ensuring alignment across the human parts in both 2D image generation and 3D texture reconstruction. The second block of Tab. 2 shows that utilizing both the front-view image and text prompts leads to visually plausible 2D image generation and accurate texture reconstruction. The front-view image provides fine-grained human appearance details, ensuring that the reconstructed textures closely match the appearance of the input image. The text prompts provide semantic information for each part, such as cloth color and style, and contribute to inferring human textures, especially in invisible human regions. Therefore, by benefiting from PartDiffusion, which effectively aggregates information from the front-view image and text prompts, PartTexturer achieves high-quality and coherent 3D human textures.

5.4. Comparison with state-of-the-art methods

Fig. 6 and Tab. 3 show that our PARTE achieves superior texture reconstruction than both regression-based and optimization-based reconstruction methods. The regression-based methods train a network to infer 3D textures from a single image but struggle to capture the structure of human parts, as they learn global context without explicit local guidance on 3D human parts. Accordingly, their reconstructed textures often exhibit part misalignment, causing textures to bleed across human parts. The optimization-based methods, which leverage diffusion




Figure 6. Qualitative comparison with existing 3D human reconstruction methods: SiTH [24], HumanRef [82], and TeCH [30].

Methods	THuman2.1					HuMMan				
	3D geometry reconstruction		3D texture reconstruction			3D geometry reconstruction		3D texture reconstruction		
	P2S [↓]	CD [↓]	PSNR [↑]	LPIPS [↓]	Part IoU [↑]	P2S [↓]	CD [↓]	PSNR [↑]	LPIPS [↓]	Part IoU [↑]
* Regression-based methods										
PIFu [61]	3.623	3.685	19.382	0.164	0.109	3.512	3.516	19.355	0.167	0.156
2K2K [20]	3.552	3.109	20.373	0.131	0.515	3.346	3.367	20.662	0.110	0.491
SiTH [24]	3.047	3.440	20.692	0.120	0.535	2.872	3.641	20.396	0.124	0.457
* Optimization-based methods										
HumanRef [82]	3.185	3.199	21.302	0.113	0.576	3.312	3.591	20.974	0.101	0.442
SIFU [85]	3.042	3.021	21.491	0.108	0.471	2.808	3.027	21.578	0.098	0.437
TeCH [30]	2.984	3.008	21.089	0.108	0.588	2.718	2.667	21.642	0.097	0.437
PARTE(Ours)	2.984	3.008	22.175	0.096	0.641	2.718	2.667	22.076	0.083	0.589

Table 3. Quantitative comparisons with existing 3D human reconstruction methods.

networks for high-quality texture reconstruction, also suffer from part misalignment. Their diffusion networks heavily rely on text prompts, which describe the overall appearance of the image without distinguishing between different human parts. Such reliance on text prompts can lead to unintended texturing, as the information from the text prompt is mapped onto incorrect human parts. Unlike these methods, our framework explicitly incorporates part segmentation from the input image, ensuring that textures are accurately assigned to their corresponding human parts, enhancing the part alignment. Additionally, our framework decomposes the texture reconstruction problem into two specialized tasks: part segmentation and part texturing, resulting in high-quality part segmentation and texturing results. Without explicit part information, the framework is required to infer both the part structure and the overall human texture, making the reconstruction problem significantly more chal-

lenging. In this regard, our framework divides the complex reconstruction process into two stages, PartSegmenter and PartTexturer, each dedicated to the human part structure and human part texture, respectively. By specializing both stages, our framework enhances the performance of both part segmentation and texture reconstruction, leading to highly detailed and realistic 3D human reconstructions.

6. Conclusion

We present **PARTE**, a framework that reconstructs a high-quality 3D human from a single image by explicitly utilizing 3D human part information for reconstructing human texture. Our framework effectively mitigates texture misalignment across human parts by leveraging our proposed PartSegmenter for 3D part segmentation and PartTexturer for part-guided texturing. As a result, PARTE achieves state-of-the-art performance in 3D human reconstruction.

Acknowledgements. This work was supported in part by the IITP grants [No.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University), No.2021-0-02068, No.2023-0-00156, IITP-2025-RS-2020-II201819, RS-2025-02653113, and RS-2025-25441838]. This work was also supported by the Industrial Technology Alchemist Project [No. RS-2024-00432410] and the Technology Innovation Program [RS-2025-02653087], both funded by MOTIE, Korea.

References

- [1] Rameen Abdal, Wang Yifan, Zifan Shi, Yinghao Xu, Ryan Po, Zhengfei Kuang, Qifeng Chen, Dit-Yan Yeung, and Gordon Wetzstein. Gaussian shell maps for efficient 3D human generation. In *CVPR*, 2024. 3
- [2] Badour AlBahar, Shunsuke Saito, Hung-Yu Tseng, Changil Kim, Johannes Kopf, and Jia-Bin Huang. Single-image 3D human digitization with shape-guided diffusion. In *SIG-GRAPH Asia*, 2023. 1
- [3] Thimo 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
- [4] Thimo Alldieck, Hongyi Xu, and Cristian Sminchisescu. imGHUM: Implicit generative models of 3D human shape and articulated pose. In *ICCV*, 2021. 3
- [5] Thimo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3D reconstruction of humans wearing clothing. In *CVPR*, 2022. 2
- [6] Thimo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3D reconstruction of humans wearing clothing. In *CVPR*, 2022. 1
- [7] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. MultiDiffusion: Fusing diffusion paths for controlled image generation. *Proceedings of Machine Learning Research*, 2023. 3
- [8] Zhongang Cai, Daxuan Ren, Ailing Zeng, Zhengyu Lin, Tao Yu, Wenjia Wang, Xiangyu Fan, Yang Gao, Yifan Yu, Liang Pan, Fangzhou Hong, Mingyuan Zhang, Chen Change Loy, Lei Yang, and Ziwei Liu. HuMMan: Multi-modal 4D human dataset for versatile sensing and modeling. In *ECCV*, 2022. 1, 6
- [9] Yukang Cao, Yan-Pei Cao, Kai Han, Ying Shan, and Kwan-Yee K Wong. DreamAvatar: Text-and-shape guided 3D human avatar generation via diffusion models. In *CVPR*, 2024. 3
- [10] Jinnan Chen, Chen Li, Jianfeng Zhang, Lingting Zhu, Buzhen Huang, Hanlin Chen, and Gim Hee Lee. Generalizable human gaussians from single-view image. *arXiv preprint arXiv:2406.06050*, 2024. 1
- [11] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3D: Disentangling geometry and appearance for high-quality text-to-3D content creation. In *ICCV*, 2023. 2
- [12] Zhaoxi Chen, Fangzhou Hong, Haiyi Mei, Guangcong Wang, Lei Yang, and Ziwei Liu. PrimDiffusion: Volumetric primitives diffusion for 3D human generation. In *NeurIPS*, 2023. 3
- [13] Enric Corona, Mihai Zanfir, Thimo Alldieck, Eduard Gabriel Bazavan, Andrei Zanfir, and Cristian Sminchisescu. Structured 3D features for reconstructing controllable avatars. In *CVPR*, 2023. 2
- [14] Haoye Dong, Xiaodan Liang, Ke Gong, Hanjiang Lai, Jia Zhu, and Jian Yin. Soft-gated warping-GAN for pose-guided person image synthesis. *NIPS*, 2018. 3
- [15] Zijian Dong, Xu Chen, Jinlong Yang, Michael J Black, Otmar Hilliges, and Andreas Geiger. AG3D: Learning to generate 3D avatars from 2D image collections. In *ICCV*, 2023. 3
- [16] Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen Change Loy, Wayne Wu, and Ziwei Liu. StyleGAN-Human: A data-centric odyssey of human generation. In *ECCV*, 2022. 2, 4, 6
- [17] Tsu-Jui Fu, Wenhan Xiong, Yixin Nie, Jingyu Liu, Barlas Oguz, and William Yang Wang. Text-guided 3D human generation from 2D collections. In *EMNLP*, 2023. 3
- [18] Xiangjun Gao, Xiaoyu Li, Chaopeng Zhang, Qi Zhang, Yanpei Cao, Ying Shan, and Long Quan. ConTex-Human: Free-view rendering of human from a single image with texture-consistent synthesis. In *CVPR*, 2024. 2
- [19] Artur Grigorev, Karim Isakov, Anastasia Ianina, Renat Bashirov, Ilya Zakharkin, Alexander Vakhitov, and Victor Lempitsky. StylePeople: A generative model of fullbody human avatars. In *CVPR*, 2021. 3
- [20] Sang-Hun Han, Min-Gyu Park, Ju Hong Yoon, Ju-Mi Kang, Young-Jae Park, and Hae-Gon Jeon. High-fidelity 3D human digitization from single 2k resolution images. In *CVPR*, 2023. 1, 2, 8
- [21] Xintong Han, Xiaojun Hu, Weilin Huang, and Matthew R Scott. ClothFlow: A flow-based model for clothed person generation. In *ICCV*, 2019. 3
- [22] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. ARCH++: Animation-ready clothed human reconstruction revisited. In *ICCV*, 2021. 1, 2
- [23] Hsuan-I Ho, Lixin Xue, Jie Song, and Otmar Hilliges. Learning locally editable virtual humans. In *CVPR*, 2023. 1
- [24] I Ho, Jie Song, Otmar Hilliges, et al. SiTH: Single-view textured human reconstruction with image-conditioned diffusion. In *CVPR*, 2024. 1, 2, 8
- [25] Fangzhou Hong, Mingyuan Zhang, Liang Pan, Zhongang Cai, Lei Yang, and Ziwei Liu. AvatarCLIP: zero-shot text-driven generation and animation of 3D avatars. *TOG*, 2022. 3
- [26] Fangzhou Hong, Zhaoxi Chen, LAN Yushi, Liang Pan, and Ziwei Liu. EVA3D: Compositional 3D human generation from 2D image collections. In *ICLR*, 2023.
- [27] Tao Hu, Fangzhou Hong, and Ziwei Liu. StructLDM: Structured latent diffusion for 3D human generation. In *ECCV*, 2024.
- [28] Xin Huang, Ruizhi Shao, Qi Zhang, Hongwen Zhang, Ying Feng, Yebin Liu, and Qing Wang. HumanNorm: Learning normal diffusion model for high-quality and realistic 3D human generation. In *CVPR*, 2024.

- [29] Yukun Huang, Jianan Wang, Ailing Zeng, He Cao, Xianbiao Qi, Yukai Shi, Zheng-Jun Zha, and Lei Zhang. DreamWaltz: Make a scene with complex 3D animatable avatars. In *NeurIPS*, 2024. 3
- [30] Yangyi Huang, Hongwei Yi, Yuliang Xiu, Tingting Liao, Ji-xiang Tang, Deng Cai, and Justus Thies. TeCH: Text-guided reconstruction of lifelike clothed humans. In *3DV*, 2024. 1, 2, 3, 8
- [31] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. ARCH: Animatable reconstruction of clothed humans. In *CVPR*, 2020. 1, 2
- [32] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *CVPR*, 2022. 2
- [33] Ruixiang Jiang, Can Wang, Jingbo Zhang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. AvatarCraft: Transforming text into neural human avatars with parameterized shape and pose control. In *ICCV*, 2023. 3
- [34] Rawal Khirodkar, Timur Bagautdinov, Julieta Martinez, Su Zhaoen, Austin James, Peter Selednik, Stuart Anderson, and Shunsuke Saito. Sapiens: Foundation for human vision models. In *ECCV*, 2024. 4, 5, 6
- [35] Byungjun Kim, Patrick Kwon, Kwangho Lee, Myunggi Lee, Sookwan Han, Daesik Kim, and Hanbyul Joo. Chupa: Carving 3D clothed humans from skinned shape priors using 2D diffusion probabilistic models. In *ICCV*, 2023. 3
- [36] Nikos Kolotouros, Thiemo Alldieck, Andrei Zanfir, Eduard Bazavan, Mihai Fieraru, and Cristian Sminchisescu. DreamHuman: Animatable 3D avatars from text. In *NeurIPS*, 2023.
- [37] Nikos Kolotouros, Thiemo Alldieck, Enric Corona, Eduard Gabriel Bazavan, and Cristian Sminchisescu. Instant 3D human avatar generation using image diffusion models. In *ECCV*, 2024. 1, 3
- [38] Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 360-degree textures of people in clothing from a single image. In *3DV*, 2019. 2
- [39] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 5
- [40] Wenbo Li, Pengchuan Zhang, Lei Zhang, Qiuyuan Huang, Xiaodong He, Siwei Lyu, and Jianfeng Gao. Object-driven text-to-image synthesis via adversarial training. In *CVPR*, 2019. 3
- [41] Yuheng Li, Yijun Li, Jingwan Lu, Eli Shechtman, Yong Jae Lee, and Krishna Kumar Singh. Collaging class-specific GANs for semantic image synthesis. In *ICCV*, 2021. 3
- [42] Tingting Liao, Xiaomei Zhang, Yuliang Xiu, Hongwei Yi, Xudong Liu, Guo-Jun Qi, Yong Zhang, Xuan Wang, Xianguyu Zhu, and Zhen Lei. High-fidelity clothed avatar reconstruction from a single image. In *CVPR*, 2023. 2
- [43] Tingting Liao, Hongwei Yi, Yuliang Xiu, Ji-xiang Tang, Yangyi Huang, Justus Thies, and Michael J. Black. TADA! text to animatable digital avatars. In *3DV*, 2024. 3
- [44] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 2, 4
- [45] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. DeepFashion: Powering robust clothes recognition and retrieval with rich annotations. In *CVPR*, 2016. 2, 4
- [46] Zhengyao Lv, Xiaoming Li, Xin Li, Fu Li, Tianwei Lin, Dongliang He, and Wangmeng Zuo. Learning semantic person image generation by region-adaptive normalization. In *CVPR*, 2021. 3
- [47] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J Black. Learning to dress 3D people in generative clothing. In *CVPR*, 2020. 1
- [48] Yifang Men, Biwen Lei, Yuan Yao, Miaomiao Cui, Zhouhui Lian, and Xuansong Xie. En3D: An enhanced generative model for sculpting 3D humans from 2D synthetic data. In *CVPR*, 2024. 1, 3
- [49] Gyeongsik Moon, Hyeongjin Nam, Takaaki Shiratori, and Kyoung Mu Lee. 3D clothed human reconstruction in the wild. In *ECCV*, 2022. 2
- [50] Hyeongjin Nam, Donghwan Kim, Jeongtaek Oh, and Kyoung Mu Lee. DeCloTH: Decomposable 3D cloth and human body reconstruction from a single image. In *CVPR*, 2025.
- [51] Panwang Pan, Zhuo Su, Chenguo Lin, Zhen Fan, Yongjie Zhang, Zeming Li, Tingting Shen, Yadong Mu, and Yebin Liu. HumanSplat: Generalizable single-image human gaussian splatting with structure priors. In *NeurIPS*, 2025. 2
- [52] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, 2019. 3
- [53] Hao-Yang Peng, Jia-Peng Zhang, Meng-Hao Guo, Yan-Pei Cao, and Shi-Min Hu. CharacterGen: Efficient 3D character generation from single images with multi-view pose canonicalization. *TOG*, 2024. 2
- [54] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. In *ICLR*, 2023. 2, 5
- [55] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, 2017. 4
- [56] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 2017. 4
- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 3
- [58] Renderpeople, 2018. <https://renderpeople.com/3d-people>. 1
- [59] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 5, 7
- [60] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 1, 2

- [61] 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. 1, 2, 8
- [62] 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. 1, 2
- [63] Akash Sengupta, Thiemo Alldieck, Nikos Kolotouros, Enric Corona, Andrei Zanfir, and Cristian Sminchisescu. DiffHuman: probabilistic photorealistic 3D reconstruction of humans. In *CVPR*, 2024. 2
- [64] Sijie Song, Wei Zhang, Jiaying Liu, and Tao Mei. Unsupervised person image generation with semantic parsing transformation. In *CVPR*, 2019. 3
- [65] Wei Sun and Tianfu Wu. Image synthesis from reconfigurable layout and style. In *ICCV*, 2019. 3
- [66] Ayça Takmaz, Jonas Schult, Irem Kaftan, Mertcan Akçay, Bastian Leibe, Robert Sumner, Francis Engelmann, and Siyu Tang. 3D segmentation of humans in point clouds with synthetic data. In *ICCV*, 2023. 4, 6, 7
- [67] Junying Wang, Jae Shin Yoon, Tuanfeng Y Wang, Krishna Kumar Singh, and Ulrich Neumann. Complete 3D human reconstruction from a single incomplete image. In *CVPR*, 2023. 2
- [68] Wentao Wang, Hang Ye, Fangzhou Hong, Xue Yang, Jianfu Zhang, Yizhou Wang, Ziwei Liu, and Liang Pan. GeneMAN: Generalizable single-image 3D human reconstruction from multi-source human data. In *ICLR*, 2025. 2
- [69] Xudong Wang, Trevor Darrell, Sai Saketh Rambhatla, Rohit Girdhar, and Ishan Misra. InstanceDiffusion: Instance-level control for image generation. In *CVPR*, 2024. 3, 7
- [70] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. ProlificDreamer: High-fidelity and diverse text-to-3D generation with variational score distillation. In *NeurIPS*, 2023. 2
- [71] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Photo Wake-Up: 3D character animation from a single photo. In *CVPR*, 2019. 2
- [72] Zhangyang Xiong, Di Kang, Derong Jin, Weikai Chen, Linchao Bao, Shuguang Cui, and Xiaoguang Han. Get3DHuman: Lifting StyleGAN-human into a 3D generative model using pixel-aligned reconstruction priors. In *ICCV*, 2023. 3
- [73] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. ICON: Implicit clothed humans obtained from normals. In *CVPR*, 2022. 1
- [74] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J Black. ECON: Explicit clothed humans optimized via normal integration. In *CVPR*, 2023. 1
- [75] Xiangyu Xu and Chen Change Loy. 3D human texture estimation from a single image with transformers. In *ICCV*, 2021. 2
- [76] Zhongcong Xu, Jianfeng Zhang, Jun Hao Liew, Jiashi Feng, and Mike Zheng Shou. XAGen: 3D expressive human avatars generation. In *NeurIPS*, 2023. 3
- [77] Yuxuan Xue, Xianghui Xie, Riccardo Marin, and Gerard Pons-Moll. Human-3Diffusion: Realistic avatar creation via explicit 3D consistent diffusion models. *arXiv preprint arXiv:2406.08475*, 2024. 1
- [78] Zhengyuan Yang, Jianfeng Wang, Zhe Gan, Linjie Li, Kevin Lin, Chenfei Wu, Nan Duan, Zicheng Liu, Ce Liu, Michael Zeng, et al. ReCo: Region-controlled text-to-image generation. In *CVPR*, 2023. 3
- [79] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4D: Real-time human volumetric capture from very sparse consumer RGBD sensors. In *CVPR*, 2021. 1, 6
- [80] Xiaoyu Zhan, Jianxin Yang, Yuanqi Li, Jie Guo, Yanwen Guo, and Wenping Wang. Semantic human mesh reconstruction with textures. In *CVPR*, 2024. 1, 2
- [81] Jichao Zhang, Enver Sangineto, Hao Tang, Aliaksandr Siarohin, Zhun Zhong, Nicu Sebe, and Wei Wang. 3D-aware semantic-guided generative model for human synthesis. In *ECCV*, 2022. 3
- [82] Jingbo Zhang, Xiaoyu Li, Qi Zhang, Yanpei Cao, Ying Shan, and Jing Liao. HumanRef: Single image to 3D human generation via reference-guided diffusion. In *CVPR*, 2024. 1, 2, 8
- [83] 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
- [84] Zechuan Zhang, Li Sun, Zongxin Yang, Ling Chen, and Yi Yang. Global-correlated 3d-decoupling transformer for clothed avatar reconstruction. In *NeurIPS*, 2023. 1
- [85] Zechuan Zhang, Zongxin Yang, and Yi Yang. SIFU: Side-view conditioned implicit function for real-world usable clothed human reconstruction. In *CVPR*, 2024. 1, 2, 8
- [86] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *ICCV*, 2021. 4
- [87] Guangcong Zheng, Xianpan Zhou, Xuewei Li, Zhongang Qi, Ying Shan, and Xi Li. LayoutDiffusion: Controllable diffusion model for layout-to-image generation. In *CVPR*, 2023. 3
- [88] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. PaMIR: Parametric model-conditioned implicit representation for image-based human reconstruction, 2021. 1, 2
- [89] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. SEAN: Image synthesis with semantic region-adaptive normalization. In *CVPR*, 2020. 3