

# ConTex-Human: Free-View Rendering of Human from a Single Image with Texture-Consistent Synthesis

Xiangjun Gao<sup>1</sup> Xiaoyu Li<sup>2</sup>† Chaopeng Zhang<sup>2</sup> Qi Zhang<sup>2</sup> Yanpei Cao<sup>2</sup>  
 Ying Shan<sup>2</sup> Long Quan<sup>1</sup>

<sup>1</sup>The Hong Kong University of Science and Technology <sup>2</sup>Tencent AI Lab

† Corresponding author




Figure 1. Our approach “ConTex-Human” can achieve texture-consistent free-view human rendering with high-fidelity using only a single image on different datasets. The left two are from *SSHQ*, right two are from *THuman2.0*. (Check [project page](#) for more visual results.)

## Abstract

In this work, we propose a method to address the challenge of rendering a 3D human from a single image in a free-view manner. Some existing approaches could achieve this by using generalizable pixel-aligned implicit fields to reconstruct a textured mesh of a human or by employing a 2D diffusion model as guidance with the Score Distillation Sampling (SDS) method, to lift the 2D image into 3D space. However, a generalizable implicit field often results in an over-smooth texture field, while the SDS method tends to lead to a texture-inconsistent novel view with the input image. In this paper, we introduce a texture-consistent back view synthesis module that could transfer the reference image content to the back view through depth and text-guided attention injection. Moreover, to alleviate the color distortion that occurs in the side region, we propose a visibility-aware patch consistency regularization for texture map-

ping and refinement combined with the synthesized back view texture. With the above techniques, we can achieve high-fidelity and texture-consistent human rendering from a single image. Experiments conducted on both real and synthetic data demonstrate the effectiveness of our method and show that our approach outperforms previous baseline methods.

## 1. Introduction

Free-view human synthesis or rendering is essential for various applications, including virtual reality, electronic games, and movie production. Traditional approaches often require a dense camera rig or depth sensors [5, 9] to reconstruct the geometry and refine the texture of the subject, resulting in a tedious and time-consuming process.

Recently, with the advent of implicit fields, remarkable progress has been made in 3D human free-view synthesis from a single RGB image. Several methods [43, 44, 57] ac-

comply with this by using pixel-aligned 2D image features as input conditions for subsequent occupancy and color predictions of corresponding 3D points. These methods extract a mesh from the implicit field and predict the vertex colors for free-view rendering. Other works [11, 14, 20] construct generalizable human neural radiance fields trained on multi-view images, enabling the recovery of 3D humans from a single input image during testing. However, both of these methods tend to produce over-smooth and less fine details due to the smoothness bias of implicit fields and the challenge of inferring the geometry and texture of the entire body from just a single input. To forecast the invisible areas, some methods [28, 39, 48] incorporate a 2D text-to-image diffusion model as guidance and conduct score-distillation-sampling (SDS) to optimize a 3D representation from a single image. This approach can also be applied to human images, as shown in our concurrent work TeCH [16]. Nevertheless, the SDS methods tend to produce over-saturation predictions and may not achieve texture-consistent results with the input reference image in the invisible areas, especially for back view images, even when an accurate text prompt is given.

In this paper, we aim to achieve high-fidelity, texture-consistent human free-view rendering using only a single input image, as shown in Figure 1, which presents significant challenges. We propose an innovative framework named “ConTex-Human”. Under this framework, we decompose our ultimate goal into two key sub-targets. The first one involves generating a texture-consistent back view with fine details. The second sub-target is to paint the side and invisible region with reasonable texture after mapping the input reference and back view onto the reconstructed geometry.

For back view synthesis, we draw inspiration from recent 2D image/video editing methods [3, 29, 38, 50], which are capable of preserving the style and texture of the original image during the editing process. Our key idea is to query image content from the input reference image to generate a texture-consistent human back view through attention injection, guided by text prompts. However, naively generating the back view using only text prompts would lead to misalignment between the back view image and human geometry. Therefore, we control this process with the depth map as guidance to ensure that the generated back view layout is well-aligned with human geometry.

In addition to mapping the reference and back view images onto the geometry representation during the optimization, we also need to paint the side region and invisible region. An intuitive solution would be to perform the SDS method on the given person, leveraging the 2D Diffusion model prior. However, only using the SDS loss results in color distortion and over-saturation of texture. To address this problem, we propose a *visibility-aware patch consistency*

*loss* that mitigates the inconsistent side view texture. This approach ensures that the pixel values in the side and invisible regions are close to their neighboring pixels in the front or back regions.

We evaluate our approach on both the synthetic dataset THuman2.0 which has 3D textured scans as ground truth and the real dataset SSHQ which includes people in various poses, clothing, and shapes. Across the experiments, our approach exhibits significant performance in both quantitative and qualitative comparisons.

In summary, our contributions are listed as follows:

- We present an innovative framework called “ConTex-Human”, which can achieve high-fidelity free-view human rendering with consistent texture using single image.
- We design a depth and text prompt conditioned back view synthesis module that can maintain texture style and details consistent with the reference image
- We proposed a texture mapping and refinement module with a visibility-aware patch consistency loss to synthesize the consistent pixels in invisible areas.

## 2. Related Work

### 2.1. Single Image Human Recon. and Rendering

Reconstructing and Rendering 3D humans from a single image is an ill-conditioned problem that necessitates inferring the geometry or even appearance of the whole body with only one observation. Therefore, a strong prior is usually required to address this issue. Traditionally, parametric body models such as SMPL [26] are used to estimate the shape and pose from a single image [2, 12, 19, 22, 34, 36]. However, the SMPL mesh representation fails to model complex topologies like dresses and hair. Recently, implicit representations such as NeRF [32], occupancy [30], and SDF [35] have shown impressive results with the ability to model arbitrary topologies and are used to model 3D clothed humans [1, 11, 17, 23, 43, 44, 51, 52, 57]. Some of these works focus on the reconstruction of geometry [6, 44, 51, 52], while PIFu [43], PAMIR [57], ARCH [17], PHORHUM [1] and S3F [6] would also predict the texture from the image for novel view synthesis. A common limitation of these methods is the requirement of large-scale accurate 3D textured scans as training data. They also tend to predict blurry and over-smooth texture. Recently, with the emergence of large pre-trained models like CLIP [40] and Stable Diffusion [41], these model priors are also utilized to guide human reconstruction. ELICIT [15] utilizes the 3D body shape geometry prior and the visual clothing prior with the CLIP models to create plausible content in the invisible areas of animatable avatar. And TeCH [16] gives the descriptive prompts to the personalized text-to-image diffusion model to learn the invisible appearance through Score Distillation Sampling. Due to the lim-

ited expressive ability of text prompts, TeCH suffers from inconsistent texture in the generated areas. In this work, we can reconstruct texture-consistent 3D humans with the help of our texture-consistent back view synthesis method.

## 2.2. Image-to-3D Generation

Recent text-to-image synthesis has achieved high-fidelity generation results benefiting from diffusion models [13, 46] and large aligned image-text datasets. Based on the pre-trained 2D diffusion model, DreamFusion [37] proposes a Score Distillation Sampling (SDS) method that replaces CLIP models in Dream Fields [18] with the SDS loss to distill 3D models from the text prompt. After that, SDS is widely used in later text-to-3D works [4, 24, 31, 49, 56, 58]. Inspired by text-to-3D works, image-to-3D that reconstructs 3D models from a single image has also been explored [8, 28, 39, 48, 54]. Specifically, NeuralLift-360 [54] derives a prior distillation loss for CLIP-guided diffusion prior to lift a single image to a 3D object. RealFusion [28] and NeRDi [8] optimize the NeRF representations by minimizing a diffusion loss on novel view renderings with a pre-trained image diffusion model conditioned by a token inverted from input image. Recently, Make-it-3D [48] employs a two-stage optimization pipeline that builds textured point clouds to enhance the texture in fine stage, yielding high-quality 3D models according to the given image. Magic123 [39] proposes to use Stable Diffusion [41] as the 2D prior and viewpoint-conditioned diffusion model Zero-1-to-3 [25] as the 3D prior simultaneously for SDS loss to generate 3D content from a given image. However, these methods can only address objects with simple geometric structures and textures. For models with more complex geometry and texture, such as 3D humans, obvious artifacts and inconsistent texture would emerge in invisible parts.

## 3. Preliminary

Recently, DreamFusion [37] has revolutionized the text-to-3D field by proposing the score distillation sampling (SDS) method, which shows that 2D text-to-image diffusion models like Imagen [42] and StableDiffusion [41] can be lifted to generate 3D objects base on the text prompts without 3D data. It uses the 2D diffusion model to guide the text-to-3D generation process by optimizing a neural radiance field with a SDS loss that is described as follows:

$$\nabla_{\theta} \mathcal{L}_{sds}^{sd} = \mathbb{E}_{t, \epsilon} \left[ w(t) (\hat{\epsilon}_{\phi}(\mathbf{x}_t; y, t) - \epsilon) \frac{\partial \mathbf{x}}{\partial \theta} \right] \quad (1)$$

where  $y$  is the given text prompts,  $\mathbf{x}$  is the rendered image from the 3D representation,  $\mathbf{x}_t$  is the noisy latent after adding Gaussian noise  $\epsilon$  to  $\mathbf{x}$ ,  $\hat{\epsilon}_{\phi}$  is the predicted noise,  $w(t)$  is the weight function of different noise levels,  $\theta$  is the parameters of 3D representations, which is NeRF in DreamFusion.

The SDS loss is not only widely used in text-to-3D tasks but also garners significant attention in image-to-3D applications. To learn rich 3D priors from large-scale 3D datasets, such as Objaverse [7], Zero-1-to-3 [25] trains a viewpoint-conditioned diffusion model capable of synthesizing novel views in a feed-forward manner. Given a single image and a target camera pose  $\{\mathbf{R}, \mathbf{T}\}$  as input, Zero-1-to-3 can synthesize the corresponding novel view according to the given viewpoint. Furthermore, it can also be employed to guide the optimization of the image-to-3D process using a modified version of the SDS loss, formulated as:

$$\nabla_{\theta} \mathcal{L}_{sds}^{z123} = \mathbb{E}_{t, \epsilon} \left[ w(t) (\hat{\epsilon}_{\phi}(\mathbf{x}_t; \mathbf{I}^r, t, \mathbf{R}, \mathbf{T}) - \epsilon) \frac{\partial \mathbf{I}}{\partial \theta} \right] \quad (2)$$

where  $\mathbf{I}^r$  is reference image,  $\mathbf{I}$  is rendered novel view. In this work, we use both of them in different stages.

## 4. Method

Given a single RGB image of a human, our objective is to reconstruct the 3D representation that could render high-fidelity images of the human from various viewpoints. Our approach only requires an RGB image and its foreground mask, which can be easily obtained using an off-the-shelf background removal tool. Our method is illustrated in Figure 2 and consists of three main stages. We first employ the 2D diffusion model to lift the input human image to a radiance field in the coarse stage (Section 4.1). Next, we introduce a depth and text-guided attention injection module from the reference to synthesize a texture-consistent image in back view (Section 4.2), serving as essential information for the subsequent stage. Finally, we propose a visibility-aware patch consistency loss to reconstruct a 3D mesh for high-quality rendering in the fine stage (Section 4.3).

### 4.1. Coarse Stage: Radiance Field Reconstruction

Recent image-to-3D generation methods [28, 39, 48] that lift a single image into a 3D object often adopt Stable Diffusion (SD) as the diffusion prior. However, we found that the SD guidance frequently results in a tedious optimization process and most importantly, would lead to inconsistent multi-head problems in the optimized 3D object due to data bias in training data of the diffusion model. In the context of 3D humans, this issue can even generate multi-arm and multi-leg geometries for simple poses, let alone handling complex and diverse human poses. To address this problem, we employ the SDS loss based on the Zero-1-to-3 model as the diffusion prior as shown in Eq.2 to optimize an Instant-NGP [33] representation.

To optimize the 3D representation, We first employ  $\mathcal{L}_{mask}$  between the mask  $\tilde{\mathbf{M}}_r$  extracted from the input image and the rendered mask  $\tilde{\mathbf{M}}_r$  in the front view to restrain the human area in 3D space:

$$\mathcal{L}_{mask} = \|\mathbf{M}_r - \tilde{\mathbf{M}}_r\|_1 \quad (3)$$

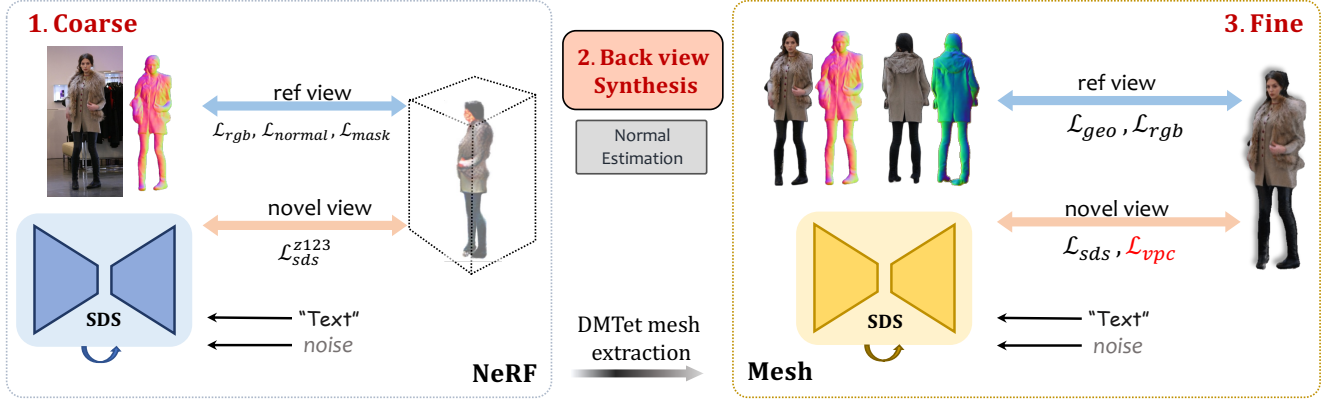


Figure 2. **Overview “ConTex-Human”.** Our framework is composed of three main stages. (1)*Coarse Stage*. Given a human image as reference, we leverage view-aware 2D diffusion model Zero123 to conduct *Score Distillation Sampling (SDS)* to optimize a NeRF. Reference view RGB and normal supervision are also added. (2)*Back view Synthesis Stage*. Coarse image and depth map are utilized to generate texture-consistent and high-fidelity back view. (3)*Fine Stage*. We convert NeRF to DMtet Mesh and optimize mesh with front/back normal map. Texture field is optimized with front/back image, Zero123/Stable-Diffusion SDS, and visibility-aware patch consistency loss  $\mathcal{L}_{vpc}$ .

In front view, RGB loss is calculated to penalize the difference between the input image  $\mathbf{I}_r$  and rendered results  $\tilde{\mathbf{I}}_r$ :

$$\mathcal{L}_{rgb} = \|\mathbf{I}_r \odot \mathbf{M}_r - \tilde{\mathbf{I}}_r \odot \mathbf{M}_r\|_1 \quad (4)$$

In addition, to enforce better geometry and accelerate the training process during optimization, we also incorporate a reference normal constraint for the normal map rendered in the front view. The reference normal is estimated using the normal estimator proposed in ECON [53]. Therefore, the normal loss between the reference normal  $\mathcal{N}_r$  and rendered normal  $\tilde{\mathcal{N}}_r$  is formulated as:

$$\mathcal{L}_{normal} = \|\mathcal{N}_r \odot \mathbf{M}_r - \tilde{\mathcal{N}}_r \odot \mathbf{M}_r\|_1 \quad (5)$$

The overall loss for the coarse stage can be formulated as a combination of  $\mathcal{L}_{sds}^{z123}$ ,  $\mathcal{L}_{mask}$ ,  $\mathcal{L}_{rgb}$  and  $\mathcal{L}_{normal}$ .

## 4.2. Texture-Consistent Back View Synthesis

Although current image-to-3D methods can generate plausible results for invisible areas for the input image, the results tend to be over-saturation, over-smooth, style-inconsistent, and low quality due to the lack of awareness of the input image during synthesizing other areas. Inspired by the recent 2D image editing methods that could maintain content and details of source images during the synthesis and editing process. Our key idea is to query image contents from the input reference image  $\mathbf{I}_r$  and integrate them to synthesize the back view image  $\mathbf{I}_b$  while maintaining the consistent texture details, this process is guided by the text prompt  $\mathbf{T}$  and depth map  $\mathbf{D}$ .

Depth map  $\mathbf{D}$  is able to guide the layout of the  $\mathbf{I}_b$ , which is essential for the fine stage to map the texture onto the geometry seamlessly. Text prompt  $\mathbf{T}$  depicts the human information style such as gender, hair color and style, clothing color and type, etc. Based on the guidance, we propose

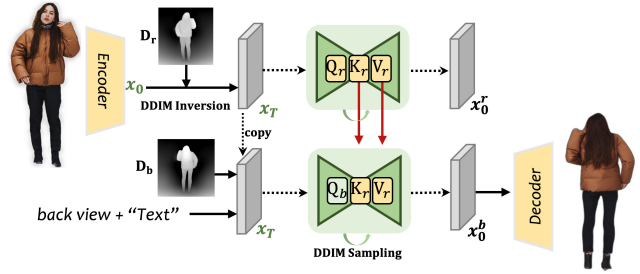


Figure 3. **Illustration of texture-consistent back view synthesis.** We firstly encode the reference image through *SD-Depth* encoder to latent code  $x_0$ . Then DDIM inversion [47] sampling is conducted to get start code  $x_T$ . The back view is synthesized from  $x_T$  through our attention injection method, which is modulated by the depth map.

a depth and text-conditioned texture-consistent back view synthesis module, which utilizes the pre-trained *Depth-Conditioned Stable Diffusion* model and synthesizes a much more highly detailed back view image than the previous methods.

Our texture-consistent back view synthesis module is shown in Figure 3. We firstly encode the original front image  $\mathbf{I}_r$  through *SD* encoder to a latent code  $x_0$ . Then the DDIM inversion [47] sampling is conducted on  $x_0$  which is concatenated with the front view depth  $\mathbf{D}_r$  iteratively to get the start noise latent code  $x_T$ . For the back view synthesis,  $x_T$  is copied as the start noise latent code of the back view image, which is concatenated with the back view depth  $\mathbf{D}_b$  for the subsequent DDIM sampling. Utilizing  $\mathbf{D}_b$  as conditional information to control the layout, the generated back view is well-aligned with both the coarse stage NeRF and the fine stage Mesh. Moreover, we incorporate the phrase “back view” into the original text prompt  $\mathbf{T}$  to

guide the back view synthesis through cross-attention in the SD model.

In addition to the back view synthesis, the front view synthesis process from  $\mathbf{x}_T$  and  $\mathbf{D}_T$  is also conducted simultaneously. During the back view and front view synthesis, for the specific time step  $t$ , we employ an attention injection method to transfer the key feature  $\mathbf{K}_r$  and value feature  $\mathbf{V}_r$  in attention layers from the front view branch to the back view branch. In the meanwhile, the back view branch maintains its original query feature  $\mathbf{Q}_b$  in attention layers. The attention feature transfer is performed iteratively to synthesize the back view. With these proposed operations, the detailed texture from the front view image can be transferred to the back view, simultaneously, maintaining the back view depth layout that is view-consistent with the front view geometry and being well aligned in accordance with the original text description.

### 4.3. Fine Stage: High-Fidelity Mesh Reconstruction

The coarse stage generates only a rough geometry and low-quality texture, represented by a density field and a color field. Therefore, we introduce a fine stage to refine the geometry and texture from the coarse stage by utilizing the content details in the reference image and generated the back view image from our method. Compared to recent works [16, 39, 48] that employ the SDS loss to optimize the full texture which could suffer from over-saturation, blurriness, over-smoothing, and style inconsistency in the generated areas. Our solution, which maps the generated back view that is style and texture-consistent with the front view onto the refined geometry, combining the proposed visibility-aware patch consistency regularization, achieves more 3D-consistent and high-fidelity results.

#### 4.3.1 Geometry Reconstruction

We adopt DM Tet [45], a hybrid SDF-Mesh representation, for the reconstruction in the fine stage, which is capable of generating high-resolution 3D shapes and allows for efficient differentiable rendering. To initialize DM Tet, we set the SDF value of each vertex  $v_i$  using the density field from the coarse stage, and the deformation vector  $\Delta v_i$  is set to  $\mathbf{0}$ . During geometry optimization, a triangle mesh is extracted from DM Tet. We employ a differential rasterizer [21] to render the normal map from a given viewpoint.

To regularize the geometry during optimization, we also employ a normal constraint as in the coarse stage. One straightforward way is to estimate the normal maps of both the front and back view from the reference image as the supervision using an existing normal estimator in ECON [53]. However, alignment issues arise between the estimated back view normal and the reconstructed geometry due to the different camera settings. In view of this, we alternatively

adopt the estimated back normal  $\mathcal{N}_b$  from the synthesized texture-consistent back view that is generated in Section 4.2. This normal is well-aligned with our initialized geometry and synthesized back view.

Given that the reference view normal and back view normal encompass most of the human region, a reasonable transition between the reference and back views can be achieved after applying mesh normal smoothness and laplacian smoothness constraints. Finally, the geometry reconstruction loss in the fine stage can be written as follows:

$$\mathcal{L}_{geo} = \|\mathcal{N}_r - \tilde{\mathcal{N}}_r\|_2 + \|\mathcal{N}_b - \tilde{\mathcal{N}}_b\|_2 + \mathcal{L}_{smooth}$$

where  $\tilde{\mathcal{N}}_r$  and  $\tilde{\mathcal{N}}_b$  are the rendered mesh normal maps using Nvdiffrast [21], respectively.  $\mathcal{N}_r$  and  $\mathcal{N}_b$  are the reference and back view target normal maps, respectively, estimated using ECON [53] normal estimator for the reference image and generated back view image.

#### 4.3.2 Texture Mapping and Refinement

After the geometry reconstruction in the fine stage, the next step is to generate texture by mapping the reference front image  $\mathbf{I}_r$  and the generated back view image  $\mathbf{I}_b$  to the refined geometry. Similar to [24], we adopt an Instant-NGP to represent the 3D texture field. For each pixel  $\mathbf{x}_i$  in the sampled image, we first calculate its ray-mesh intersection’s 3D position  $\mathbf{p}_i$ . Then a latent feature is interpolated from the Instant-NGP feature grid and is fed to a tiny layer MLP network to decode a color value. The texture field is first regularized using the front reference image  $\mathbf{I}_r$  and the generated back view image  $\mathbf{I}_b$  as the supervision:

$$\mathcal{L}_{rgb} = \|\mathbf{I}_r - \tilde{\mathbf{I}}_r\|_2 + \|\mathbf{I}_b - \tilde{\mathbf{I}}_b\|_2 \quad (6)$$

where  $\tilde{\mathbf{I}}_r$  and  $\tilde{\mathbf{I}}_b$  are the rendered front and back images from the texture field, respectively.

Although the front and back view images could cover most of the texture for the human, there are still some missing textures in the side view and self-occluded region. To complete the missing texture, similar to [39], we adopt a SDS combination of both Stable-Diffusion and Zero-1-to-3 models to optimize the texture field:

$$\mathcal{L}_{sds} = \lambda_1 \mathcal{L}_{sds}^{sd} + \lambda_2 \mathcal{L}_{sds}^{z123} \quad (7)$$

The combined SDS loss can fill the missing region guided by text prompts. However, in many cases, there is an obvious texture transition and inconsistent style in the filled part between the front and back view images. To address this problem, we propose a *visibility-aware patch consistency loss* for refinement, which could alleviate the inconsistent side view texture as shown in Figure 4. To be specific, for each pixel in the front view image and back view image, we find its intersection with the corresponding mesh triangle

face through rasterization. The vertices on the face closest to the intersection are set to 1, indicating that they are visible to  $\mathbf{I}_r$  or  $\mathbf{I}_b$ . Vertices that are invisible to  $\mathbf{I}_r$  and  $\mathbf{I}_b$  are set to 0.

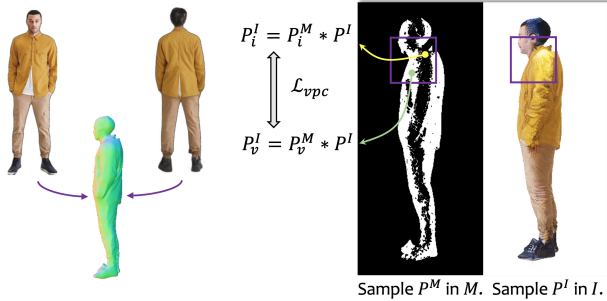


Figure 4. **Illustration of visibility-aware patch consistency loss**  $\mathcal{L}_{vpc}$ . After mapping the front/back view to the geometry, we conduct *SDS* to complete side and invisible region. To remove the color distortion caused by *SDS*, we sample random patches  $P^I$  which are divided into visible region  $P_v^I$  and invisible region  $P_i^I$ . The  $\mathcal{L}_{vpc}$  is calculated between regions.

Our key insight is that the pixels in the invisible region should have a consistent color with their neighbor visible pixels within a patch. To achieve this, we first sample a random viewpoint in camera space and render an RGB image  $I$  and its visibility map  $M$ . Then we sample a random patch  $P^I$  in  $I$  and its visibility map  $P^M$  in  $M$ . In this patch, the invisible pixels  $P_i^I$  can be calculated by ( $P_i^I = P^I * P_i^M$ ), and the visible pixels  $P_v^I$  can be calculated by ( $P_v^I = P^I * P_v^M$ ). Then we calculate the visibility-aware patch consistency loss as follows:

$$\mathcal{L}_{vpc} = \sum_{p \in P_i^I} \min_{q \in P_v^I} \|p - q\|^2 \quad (8)$$

The overall loss for the texture mapping and refinement can be formulated as a combination of  $\mathcal{L}_{rgb}$ ,  $\mathcal{L}_{sds}$  and  $\mathcal{L}_{vpc}$ . Due to space constraints, we will provide more implementation details in the *supplementary materials*.

## 5. Experiments

### 5.1. Datasets

We describe the human datasets used in the experiments, including a synthetic dataset rendered by 3D textured scans THuman2.0 [55], and a real dataset with high-quality full-body human images, SSHQ [10].

**THuman2.0** is a 3D human model dataset that contains 500 high-quality human scans captured by a dense DSLR rig. For each scan, it provides a 3D model along with the corresponding texture map. In all 500 static scans, the same person might appear multiple times with different poses and clothes. As our method is person-specific, evaluating all the data would be a tedious process. Therefore, we selected

30 subjects with different identities, poses, and clothes for evaluation. For each subject, we used the front view as the input for our method. To evaluate our method, we rendered ten views that surround the center human as ground truth novel views using PyTorch3D. We rendered images at a resolution of 648×648 pixels, where the height of the human region comprises approximately 70% of the image, resulting in a height of around 455 pixels.

**SSHQ** is a dataset consisting of high-quality full-body human images at a resolution of 1024×512. SSHQ covers a wide range of races, clothing styles, and poses. Similar to THuman2.0, we selected 30 subjects with only a single image for evaluation. We remove the background, resize the image to a resolution of 648×648 pixels, and reposition the human to occupy approximately 70% of the image’s height, ensuring it remains centered.

### 5.2. Metrics and Methods for Comparison

To evaluate our methods, we compare our method with PIFu [43] and PaMIR [57] which are single image human reconstruction methods including the texture that can be inferred in a feed-forward manner. Additionally, we also compare our results with a recent image-to-3D generation method Magic123 [39] which is a per-subject optimization method like ours. For the THuman2.0 dataset, since both the input view and novel views have the ground truth, we evaluate all methods on the rendered images in input view and novel views using commonly used metrics: PSNR, SSIM, LPIPS, and CLIP. For the SSHQ dataset that only has the input image for evaluation, following the experiments in Magic123, we adopt CLIP similarity to measure the consistency between the input image and the rendered novel view images and use LPIPS to measure the accuracy between the input image and the rendered reference view.

For TeCH [16], a concurrent work to our method, released their code two weeks before the paper submission. However, it still exhibits inconsistencies in the generated areas. We provide a comparison with it in our *suppl*.

### 5.3. Evaluation

**Results on the THuman2.0 dataset.** Table 1 summarizes the quantitative comparison of our proposed method with the baseline methods on THuman2.0 dataset. PIFu and PaMIR sample dense 3D points in 3D space and the sampling points are projected onto the input image plane to get the interpolated image features for occupancy prediction. Then an explicit mesh is extracted using the Marching-Cube[27] algorithm. For rendering, they predict a color value for each mesh vertex. Similar to us, Magic123 is a per-subject optimization method that reconstructs the given subject by lifting the 2D diffusion models to 3D using the SDS loss. In Table 1, we demonstrate that our approach surpasses all the baseline methods with respect to PSNR,

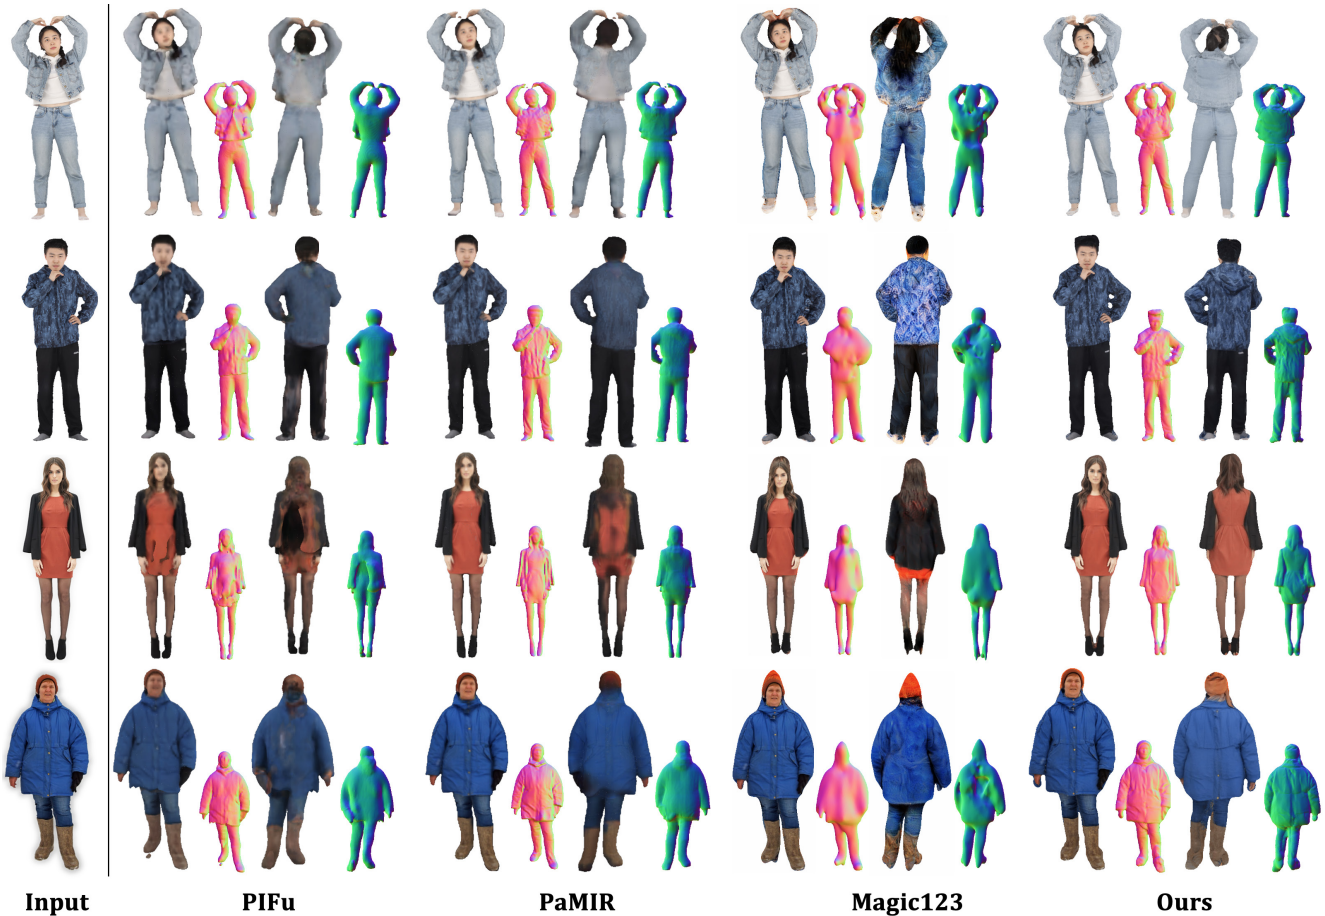


Figure 5. **Qualitative results on THuman2.0 and SSHQ dataset.** Row 1&2 are THuman2.0 samples, Row 3&4 are SSHQ samples. PIFu and PaMIR tend to predict blurred rendering results, especially in the back view. Magic123 has difficulty in predicting consistent texture. Compared with them, our methods accurately render texture-consistent and high-fidelity novel views, please *Zoom in* for the details.

LPIPS, and CLIP metrics while achieving the second-best performance in terms of SSIM. This indicates that our results are more closely aligned with the ground truth.

In row 1&2 of Figure 5, we present qualitative results of our method and baseline methods in the front view and back view rendered images. Our method produces a clearer, more detailed, and more photo-realistic front view and back view than PIFu and PaMIR. Magic123 effectively preserves the details of the input image by employing an RGB loss in the input view. However, it generates over-saturated and texture-inconsistent results in the back view due to the SDS loss. In contrast, our results can produce a more realistic result and consistent texture than Magic123 in different views. Our method can also handle complicated textures as the jacket shown in row 2 of Figure 5. For more visual results, please refer to our *supplementary materials*.

**Results on the SSHQ dataset.** In addition to evaluating the performance on the rendered dataset, we also evaluate all methods using the real images from the SSHQ dataset. As there is no ground truth for the novel view, we compute

Table 1. **Quantitative comparison of our method** with PIFu, PaMIR, Magic123 on THuman2.0 and SSHQ datasets in terms of SSIM, PSNR, LPIPS, and CLIP. ( $\uparrow$  means higher is better,  $\downarrow$  means lower is better.)

Method	THuman2.0				SSHQ	
	SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$	LPIPS $\downarrow$	CLIP $\uparrow$
PIFu	0.921	20.4	0.079	0.889	0.068	0.873
PaMIR	<b>0.925</b>	21.0	0.072	0.913	0.064	0.887
Magic123	0.903	18.8	0.099	0.910	<b>0.056</b>	0.882
Ours	0.923	<b>21.4</b>	<b>0.063</b>	<b>0.932</b>	0.059	<b>0.903</b>

the LPIPS metric in the front view and the CLIP similarity in the novel view. The LPIPS metric measures how closely the rendered front view matches the input image, while the CLIP similarity evaluates the resemblance between the rendered novel view and the input image. Table 5 shows the quantitative results of our method and the baseline methods. We are the second-best in LPIPS in the front view which is marginally lower than Magic123, but we are the best in the

Table 2. Comparison between **no back view synthesis** (w/o back), **no visibility-aware patch consistency loss** (w/o VPC), and **full model** on THuman2.0 dataset. (↑ means higher is better, ↓ means lower is better.)

Method	THuman2			
	SSIM↑	PSNR↑	LPIPS↓	CLIP↑
w/o back & w/o VPC	0.9126	20.188	0.9051	0.0778
with back & w/o VPC	0.9220	21.27	0.9328	0.0658
Full model	<b>0.9231</b>	<b>21.35</b>	<b>0.9315</b>	<b>0.0634</b>



Figure 6. **Qualitative results between no texture-consistent back view synthesis module and full model** on THuman2.0 and SSHQ.

CLIP similarity, showing that our rendered novel views are closer to the input image.

In rows 3 and 4 of Figure 5, we present a qualitative comparison, which shows that our model is capable of producing more detailed front and back views than PIFu and PaMIR, as well as much more consistent texture with the input image than Magic123. Furthermore, our model is even able to handle various clothing types, such as loose coats, boots, dresses, and hats, as shown in Figure 5. Although our primary focus is not on geometry, we demonstrate an improved back view normal map with enhanced details.

#### 5.4. Ablation Study

We conduct ablation studies to analyze how the texture-consistent back view synthesis module and the visibility-aware patch consistency loss in the texture mapping and refinement module affect the performance of our methods. These ablation studies for quantitative results are conducted on the THuman2.0 dataset, as it provides the ground truth data in novel views. For the qualitative results, we test the cases in both the THuman2.0 and SSHQ datasets.

**Texture-consistent back view synthesis.** To validate the importance of using a texture-consistent back view synthesis, we remove the synthesized back view image during the texture optimization in the fine stage, which means we remove the back view regularization in  $\mathcal{L}_{rgb}$  and the visibility-aware patch consistency loss  $\mathcal{L}_{vpc}$  that also relies on the back view image. In this setup, our method is more similar to previous image-to-3D methods. As shown in Table 2, the performance drops significantly when the texture-consistent back view image is removed, indicating that this




Figure 7. **Qualitative results between no visibility-aware patch consistency loss and full model** on THuman2.0 and SSHQ.

design is critical for the single image 3D human free-view rendering. The visual comparison is shown in Figure 6. It can be observed that, without the texture-consistent back view, the textures tend to be of significantly lower quality and, most importantly, lack consistency with the input view.

**Visibility-aware patch consistency loss.** Additionally, we validate the effectiveness of the proposed visibility-aware patch consistency loss (VPC) by excluding it from the optimization process. The performance also decreases as shown in Table 2. The visual examples presented in Figure 7 show that without the VPC loss, severe color distortion, as well as color inconsistency, would appear in the side region. We attribute this inconsistency to the inability of the SDS optimization to guide the model towards a texture that is consistent with the front view.

#### 6. Limitations

The depth and text-conditioned back view synthesis, along with visibility-aware patch consistency loss, enable us to achieve remarkable free-view human rendering with consistent texture. However, there are a few limitations. 1) We are unable to generate very impressive high-quality geometry. The resulting mesh exhibits coarse geometry in the hand and foot regions. Besides, if the coarse stage produces geometry with concave areas or significantly incorrect poses, the fine-stage mesh refinement cannot adequately compensate for these errors. 2) Although the side and invisible regions exhibit color-consistent predictions, their quality is not as high as that of the front and back views, and they occasionally contain some noise. 3) Similar to NeRF, our proposed method is trained in a person-specific setting, which requires over one hour to achieve training.

#### 7. Conclusion

In this paper, we introduced a novel framework for single image free-view 3D human rendering. We proposed a module for texture-consistent and high-fidelity back view synthesis, which is well-aligned with the input reference image. The texture mapping module with visibility-aware patch consistency loss is proposed for side and invisible region inpainting. Experiments on the THuman2.0 and SSHQ demonstrated that the proposed model achieves state-of-the-art performances on free-view image synthesis.



## References

- [1] Thimeo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3d reconstruction of humans wearing clothing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1506–1515, 2022. 2
- [2] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 561–578. Springer, 2016. 2
- [3] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiaohu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. *arXiv preprint arXiv:2304.08465*, 2023. 2
- [4] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023. 3
- [5] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (ToG)*, 34(4):1–13, 2015. 1
- [6] Enric Corona, Mihai Zanfir, Thimeo Alldieck, Eduard Gabriel Bazavan, Andrei Zanfir, and Cristian Sminchisescu. Structured 3d features for reconstructing controllable avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16954–16964, 2023. 2
- [7] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. 3
- [8] Congyue Deng, Chiyu Jiang, Charles R Qi, Xinchun Yan, Yin Zhou, Leonidas Guibas, Dragomir Anguelov, et al. Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20637–20647, 2023. 3
- [9] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4d: Real-time performance capture of challenging scenes. *ACM Transactions on Graphics (ToG)*, 35(4):1–13, 2016. 1
- [10] 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 *European Conference on Computer Vision*, pages 1–19. Springer, 2022. 6
- [11] Xiangjun Gao, Jiaolong Yang, Jongyoo Kim, Sida Peng, Zicheng Liu, and Xin Tong. Mps-nerf: Generalizable 3d human rendering from multiview images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2
- [12] Peng Guan, Alexander Weiss, Alexandru O Balan, and Michael J Black. Estimating human shape and pose from a single image. In *2009 IEEE 12th International Conference on Computer Vision*, pages 1381–1388. IEEE, 2009. 2
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [14] Shoukang Hu, Fangzhou Hong, Liang Pan, Haiyi Mei, Lei Yang, and Ziwei Liu. Sherf: Generalizable human nerf from a single image. *arXiv preprint arXiv:2303.12791*, 2023. 2
- [15] Yangyi Huang, Hongwei Yi, Weiyang Liu, Haofan Wang, Boxi Wu, Wenxiao Wang, Binbin Lin, Debing Zhang, and Deng Cai. One-shot implicit animatable avatars with model-based priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8974–8985, 2023. 2
- [16] Yangyi Huang, Hongwei Yi, Yuliang Xiu, Tingting Liao, Jiaxiang Tang, Deng Cai, and Justus Thies. Tech: Text-guided reconstruction of lifelike clothed humans. *arXiv preprint arXiv:2308.08545*, 2023. 2, 5, 6
- [17] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3093–3102, 2020. 2
- [18] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 867–876, 2022. 3
- [19] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7122–7131, 2018. 2
- [20] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. *Advances in Neural Information Processing Systems*, 34:24741–24752, 2021. 2
- [21] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics*, 39(6), 2020. 5
- [22] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6050–6059, 2017. 2
- [23] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8662–8672, 2023. 2
- [24] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler,

- Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 300–309, 2023. 3, 5
- [25] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9298–9309, 2023. 3
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866, 2023. 2
- [27] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353, 1998. 6
- [28] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8446–8455, 2023. 2, 3
- [29] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 2
- [30] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4460–4470, 2019. 2
- [31] Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12663–12673, 2023. 3
- [32] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [33] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 3
- [34] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *2018 international conference on 3D vision (3DV)*, pages 484–494. IEEE, 2018. 2
- [35] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. 2
- [36] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3d human pose and shape from a single color image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 459–468, 2018. 2
- [37] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 3
- [38] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. *arXiv preprint arXiv:2303.09535*, 2023. 2
- [39] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. 2, 3, 5, 6
- [40] 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [41] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2, 3
- [42] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding, 2022. 3
- [43] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2304–2314, 2019. 1, 2, 6
- [44] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 84–93, 2020. 1, 2
- [45] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *Advances in Neural Information Processing Systems*, 34:6087–6101, 2021. 5
- [46] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015. 3
- [47] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 4
- [48] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity

- 3d creation from a single image with diffusion prior. *arXiv preprint arXiv:2303.14184*, 2023. [2](#), [3](#), [5](#)
- [49] 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. *arXiv preprint arXiv:2305.16213*, 2023. [3](#)
- [50] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7623–7633, 2023. [2](#)
- [51] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J Black. Econ: Explicit clothed humans obtained from normals. *arXiv preprint arXiv:2212.07422*, 2022. [2](#)
- [52] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. Icon: Implicit clothed humans obtained from normals. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13286–13296. IEEE, 2022. [2](#)
- [53] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J Black. Econ: Explicit clothed humans optimized via normal integration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 512–523, 2023. [4](#), [5](#)
- [54] DeJia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4479–4489, 2023. [3](#)
- [55] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgb-d sensors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5746–5756, 2021. [6](#)
- [56] Minda Zhao, Chaoyi Zhao, Xinyue Liang, Lincheng Li, Zeng Zhao, Zhipeng Hu, Changjie Fan, and Xin Yu. Efficientdreamer: High-fidelity and robust 3d creation via orthogonal-view diffusion prior. *arXiv preprint arXiv:2308.13223*, 2023. [3](#)
- [57] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE transactions on pattern analysis and machine intelligence*, 44(6): 3170–3184, 2021. [1](#), [2](#), [6](#)
- [58] Joseph Zhu and Peiye Zhuang. Hifa: High-fidelity text-to-3d with advanced diffusion guidance. *arXiv preprint arXiv:2305.18766*, 2023. [3](#)