

## SHERF: Generalizable Human NeRF from a Single Image

Shoukang Hu<sup>1\*</sup> Fangzhou Hong<sup>1\*</sup> Liang Pan<sup>1</sup> Haiyi Mei<sup>2</sup> Lei Yang<sup>2</sup> Ziwei Liu<sup>1</sup>  
<sup>1</sup>S-Lab, Nanyang Technological University <sup>2</sup>SenseTime Research




Figure 1: **SHERF is a single image-based generalizable Human NeRF.** With just one inference pass on a single image, SHERF reconstructs Human NeRF in the canonical space which can be driven and rendered for novel view and pose synthesis.

### Abstract

Existing Human NeRF methods for reconstructing 3D humans typically rely on multiple 2D images from multi-view cameras or monocular videos captured from fixed camera views. However, in real-world scenarios, human images are often captured from random camera angles, presenting challenges for high-quality 3D human reconstruction. In this paper, we propose **SHERF**, the first **generalizable Human NeRF model** for recovering animatable 3D humans from a single input image. SHERF extracts and encodes 3D human representations in canonical space, enabling rendering and animation from free views and poses. To achieve high-fidelity novel view and pose synthesis, the encoded 3D human representations should capture both global appearance and local fine-grained textures. To this end, we propose a bank of 3D-aware hierarchical features, including global, point-level, and pixel-aligned features, to facilitate informative encoding. Global features enhance the information extracted from the single input image and complement the information missing from the partial 2D observation. Point-level features provide strong clues of 3D human structure, while pixel-aligned features preserve more fine-grained details. To effectively integrate the 3D-aware hierarchical feature bank, we design a feature fusion transformer. Extensive experiments on THuman, RenderPeople, ZJU\_MoCap, and HuMMan datasets demonstrate that SHERF achieves state-of-the-art performance, with better generalizability for novel view and pose synthesis. Our code is available at

<https://github.com/skhu101/SHERF>.

### 1. Introduction

Human NeRFs aim to recover high-quality 3D humans from 2D observations, avoiding the need to capture ground truth 3D geometry information [58, 13, 57, 78, 54, 75, 68, 30, 31, 73, 39, 23, 15, 84]. The development of Human NeRFs addresses a long-standing scientific request and has the potential to enable real-world applications *e.g.* VR/AR. By leveraging Human NeRF, we can reconstruct 3D humans directly from 2D observations, saving time and effort to collect ground truth 3D information.

Existing Human NeRF methods can be classified into two categories. The first category focuses on reconstructing 3D humans from monocular or multi-view videos [58, 13, 57, 78, 54, 75, 68, 30, 31, 73]. These methods optimize subject-specific Human NeRF, which are time-consuming and not suitable for the rapid applications of Human NeRF. To address the slow optimization process, the second category of Human NeRF methods [39, 23, 15, 84] propose to learn generalizable Human NeRF models. These methods can reconstruct Human NeRF from a few multi-view human images in a single forward pass, which largely speeds-up the process. Although these methods can achieve acceptable performance in 3D human reconstruction, they require multi-view images under well-defined camera angles, limiting their applicability in real-world scenarios where only a single image with a random angle is available. MonoNHR [15] addresses this gap by exploring novel view synthesis from a single image.

\*Equal contribution

But it cannot animate the reconstructed Human NeRF with novel poses, still limiting its applicability.

Recovering animatable 3D humans from a single human image with generalizable Human NeRF is a challenging problem due to two main challenges. The first challenge is the **missing information** from the partial observation. Existing generalizable Human NeRF [39, 23] focus too much on local feature preservation, while struggle to complement the missing information. The second challenge is reconstructing **animatable** 3D humans from a single human image. To make animatable Human NeRF from partial observations, it is necessary to complete missing appearance while also ensuring coherent understanding of 3D human structure. This poses additional challenges beyond the task of simply completing missing information.

In this work, we propose **SHERF**, the first generalizable Human NeRF based on single image inputs. We propose a hierarchical feature bank to address the challenge of information missing from the single image input. This feature bank includes global, point-level, and pixel-aligned features, which enable informative 3D human representations encoding. The hierarchical feature bank captures both the global human structure and local fine details, which are essential for high-fidelity human NeRF reconstruction. In addition, we introduce a feature fusion transformer to effectively merge features in the hierarchical feature bank. As illustrated in Fig. 1, our method can reconstruct correct colors for visible areas and provide plausible guesses for non-observable areas. The former is attributed to the fine-grained 3D-aware features that are crucial for reconstructing accurate geometry and color details, while the latter is enabled by the global features that allow color inference of invisible parts. The combination of these abilities leads to our method’s capability of generating high-quality novel views and poses.

To address the challenge of animatability, SHERF models the 3D human representation in canonical space, making it amenable to pose transformation and rendering. We use the SMPL prior [46] to transform hierarchical features extracted from the input image to the canonical space, where they are encoded to better complete missing information and acquire the human structure information.

We evaluate SHERF on several datasets including THuman [82], RenderPeople [1], ZJU\_MoCap [59] and HuMan [9]. Our results show that SHERF outperforms previous state-of-the-art generalizable Human NeRF methods in both novel view and novel pose synthesis with single images as inputs. We also conduct a detailed analysis on the effects of varying input camera views, which provides further insights into SHERF. Our main contributions are as follows:

1) To the best of our knowledge, SHERF is the first generalizable Human NeRF model to recover animatable 3D humans from a single human image. It pushes the boundaries of Human NeRF to a more general setting and bridges the gap

of applying Human NeRF in real-world scenarios.

2) With 3D-aware hierarchical features, SHERF learns both fine-grained and global features to recover texture details and complement information missing from partial observations.

3) SHERF achieves state-of-the-art performance compared with previous generalizable Human NeRF methods [39, 23] in both novel view and novel pose synthesis on four large-scale datasets.

## 2. Related Work

**Human NeRF.** NeRF [50, 69] has inspired research in 3D human reconstruction. Human NeRF can synthesize high-fidelity novel views or poses of 3D humans, given multi-view or monocular human videos. Neural Body [58] applies sparse convolutions to model the radiance volume, while others model human NeRF in the canonical space [13, 57, 78, 54, 75, 68, 30, 31, 73] using SMPL LBS weights or optimizing LBS weights with appearance. While these methods achieve impressive results, they often require time-consuming optimization and dense observations. To address this, there has been a growing interest in generalizable human NeRF [39, 23, 15, 84, 27]. These methods require fewer observations and only one forward pass. This work also aims to develop generalizable human NeRF and tackle a more challenging scenario, recovering animatable human NeRF from a single image.

**Monocular Human Reconstruction.** Statistical 3D human models [46, 56, 32, 62, 79] have enabled the reconstruction of 3D humans from monocular observations. Using these models, researchers have estimated coarse human shapes and poses [34, 36, 37, 38]. To model the complex shape of clothed humans, mesh deformation is estimated [2, 3, 4, 60, 87]. Implicit representations, such as SDF, have been used to improve geometry quality [63, 64, 25, 42, 19, 40, 8, 80]. To take advantage of both explicit and implicit representations, researchers have explored combining these representations [6, 7, 28, 26, 85, 77, 76, 5, 17] for better generalizability and reconstruction quality. In comparison, with the advantages of NeRF, we do not need 3D ground truth for training. Moreover, we reconstruct humans in the canonical space, which can be easily driven with novel poses.

**Generalizable NeRF.** NeRF requires dense calibrated views [50, 52], but recent advances have led to the development of generalizable NeRF that can work with very few or even single views. Cross-scene multi-view aggregators [72, 12, 44, 71] can synthesize novel views by learning to aggregate sparse views. Other works [22, 43, 67, 49, 24, 65, 10] encode observations to latent space and decode to NeRF. Our work focuses on generalizable Human NeRF that encodes single human images into the canonical 3D space.




Figure 2: **SHERF Framework.** To render the target image, we first cast rays and sample points in the target space. The sample points are transformed to the canonical space through inverse LBS. We then query the corresponding 3D-aware global, point-level, and pixel-aligned features. The deformed points, combined with the bank of features, are input into the feature fusion transformer and NeRF decoder to get the RGB and density, which are further used to produce the target image through volume rendering.

### 3. Our Approach

#### 3.1. Preliminary

**NeRF** [51] learns an implicit continuous function which takes as input the 3D location  $\mathbf{x}$  and viewing direction  $\mathbf{d}$  of each point and predicts the volume density  $\sigma \in [0, \infty)$  and color value  $\mathbf{c} \in [0, 1]^3$ , i.e.,  $F_{\Phi} : (\gamma(\mathbf{x}), \gamma(\mathbf{d})) \rightarrow (\mathbf{c}, \sigma)$ , where  $F_{\Phi}$  is parameterized by a multi-layer perception (MLP) network,  $\gamma$  is the positional embedding. To render the RGB color of pixels in the target view, rays are cast from the camera origin  $\mathbf{o}$  through the pixel with the direction  $\mathbf{d}$ . Based on the classical volume rendering [33], the expected color  $\hat{C}(\mathbf{r})$  of the camera ray  $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$  is computed as

$$\hat{C}(\mathbf{r}) = \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt, \quad (1)$$

where  $t_n$  and  $t_f$  denote the near and far bounds,  $T(t) = \exp(-\int_{t_n}^t \sigma(\mathbf{r}(s))ds)$  denotes the accumulated transmittance along the direction  $\mathbf{d}$  from  $t_n$  to  $t$ . In practice, the continuous integral is approximated with the quadrature rule [47] and reduced to the traditional alpha compositing. **SMPL** [46] is a parametric human model which defines  $\beta, \theta$  to control body shapes and poses. In this work, we apply the Linear Blend Skinning (LBS) algorithm of SMPL to transform points from the canonical space to target/observation spaces. Formally, a 3D point  $\mathbf{x}^c$  in the canonical space is transformed to an observation space defined by pose  $\theta$  as  $\mathbf{x}^o = \sum_{k=1}^K w_k \mathbf{G}_k(\theta, \mathbf{J})\mathbf{x}^c$ , where  $K$  is the joint number,

$\mathbf{G}_k(\theta, \mathbf{J})$  is the transformation matrix of joint  $k$ ,  $w_k$  is the blend weight. The transformation from target/observation spaces to the canonical space, namely inverse LBS, can be defined with inverse transformation matrices.

#### 3.2. Overview

The goal of SHERF is to train a generalizable Human NeRF model which can synthesize novel views and poses of 3D humans from a single image input. For the input human image, we assume the calibrated camera parameters and the human region masks are known. We also assume the corresponding SMPL parameters  $\{\theta, \beta\}$  are given.

The overall framework of SHERF is shown in Fig. 2. The input is a single human image  $\mathbf{I}^o$  and its corresponding camera parameters  $\mathbf{P}^o$  and SMPL pose  $\theta^o$  and shape  $\beta^o$ . The output is the human rendering in the target camera view  $\mathbf{P}^t$  with SMPL pose  $\theta^t$  and shape  $\beta^t$ . To render an image in the target space, we cast rays and sample points  $\mathbf{x}^t$  along the ray.  $\mathbf{x}^t$  is transformed to the canonical space  $\mathbf{x}^c$  through the inverse LBS. We then query the bank of 3D-aware hierarchical features, i.e., global feature  $\mathbf{f}_{global}$ , point-level feature  $\mathbf{f}_{point}$  and pixel-aligned feature  $\mathbf{f}_{pixel}$ , from their corresponding feature extraction modules. To efficiently integrate features from the feature bank, we further apply a feature fusion transformer to get the fused features  $\mathbf{f}_{trans}(\mathbf{x}^c)$  for  $\mathbf{x}^c$  as follows:

$$\mathbf{f}_{trans}(\mathbf{x}^c) = \text{Attn}(\mathbf{f}_{global}(\mathbf{x}^c), \mathbf{f}_{point}(\mathbf{x}^c), \mathbf{f}_{pixel}(\mathbf{x}^c)). \quad (2)$$

The point  $\mathbf{x}^c$ , along with the fused features  $\mathbf{f}_{trans}(\mathbf{x}^c)$ , are fed into the NeRF decoder to predict the density  $\sigma$  and RGB  $\mathbf{c}$ . Finally, volume rendering is performed in the target space to render the pixel colors by integrating the density and RGB values of sampled 3D points along the rays in the target space. In the following parts, we introduce details of the hierarchical feature extraction scheme and the feature fusion transformer.

### 3.3. Hierarchical Feature Extraction

The bank of 3D-aware hierarchical features comprises global, point-level and pixel-aligned features. With both global and fine-grained information learned from the bank of hierarchical features, SHERF enhances information observable from the input image and complements the information missing from the partial observation.

**Global Feature.** Capturing the global structure and overall appearance is essential for recovering Human NeRF from partial observations. As shown in previous work [14, 21, 29, 48, 53, 55, 61], compressing the whole scene into a compact latent code  $\mathbf{f}_{global}$  helps the encoding of such global information. Therefore, as shown in Fig. 2, we compress the input image into a compact latent code using a 2D encoder. To efficiently decode the 3D representation from the compact latent code, we involve the tri-plane representation [11], which plays an important role in missing information completion. The compact latent code is first mapped to a 512-dimensional style vector through the mapping network [35]. The style vector is then fed into the style-based encoders [35] to generate features which are further reshaped to the tri-plane representation to model Human NeRF in the canonical space. Finally, points transformed to the canonical space  $\mathbf{x}^c$  are projected to 3 planes through orthogonal projection to extract the 3D-aware global features  $\mathbf{f}_{global}(\mathbf{x}^c)$ .

**Point-Level Feature.** For single image Human NeRF, it is important to recover both global structure and local details from the input image, which can be bridged by an underlying explicit human model, *i.e.* SMPL. We first extract per-point features by projecting SMPL vertices to the 2D feature map of the input image, as shown in Fig. 2. In the single human image input setting, one problem with the above feature extraction process is that only half of the SMPL vertices are visible from the input view. Therefore, to make the point-level feature aware of occlusions, we only extract features for vertices visible from the current camera  $\mathbf{P}^o$ . Then we perform inverse LBS to transform the posed vertices features to the canonical space, which are then voxelized to sparse 3D volume tensors and further processed by sparse 3D convolutions [16]. From the encoded sparse 3D volume tensors, we can extract point-level features  $\mathbf{f}_{point}(\mathbf{x}^c)$  for point  $\mathbf{x}^c$ . The point-level features are aware of the 3D human structure and local texture details, which are helpful to infer textural information of Human NeRF in a more detailed level.

**Pixel-Aligned Feature.** The point-level features extraction uses SMPL prior and spatial convolution for both global and local feature enhancement. However, it may suffer from significant information loss due to the limited SMPL mesh resolution and the voxel resolution. To compensate for the fine-grained local information missing problem, we further extract the pixel-aligned features by projecting 3D deformed points  $\mathbf{x}^c$  into the input view. As shown in Fig. 2, we transform the deformed point  $\mathbf{x}^c$  to  $\mathbf{x}^o = \text{LBS}(\mathbf{x}^c; \boldsymbol{\theta}^o, \boldsymbol{\beta}^o)$  in observation space through LBS and project it to the input view so that pixel-aligned features  $\mathbf{f}_{pixel}(\mathbf{x}^c)$  can be queried, which can be formulated as

$$\mathbf{f}_{pixel}(\mathbf{x}^c) = \Pi(\mathbf{W}(\mathbf{I}^o); \text{LBS}(\mathbf{x}^c; \boldsymbol{\theta}^o, \boldsymbol{\beta}^o)), \quad (3)$$

where  $\mathbf{W}$  is the 2D feature encoder,  $\Pi(\cdot)$  denotes the 3D-to-2D projection operator. When used in the multi-view input setting, the variance of pixel-aligned features from different views can indicate whether 3D points are near the 3D surface or not. While in our single image setting, the pixel-aligned features can not encode such implicit information, especially when the 3D points  $\mathbf{x}^c$  are far from the surface. To avoid overfitting to the uninformative pixel-aligned features, we assign different weights\* to pixel-aligned features according to the distance between the corresponding 3D deformed point  $\mathbf{x}^c$  and its nearest SMPL vertex.

### 3.4. Feature Fusion Transformer

The above hierarchical features effectively encode different levels of texture and 3D structural information. However, it is not trivial to fuse these features. Intuitively, for the observable parts, we should rely more on the pixel-aligned feature for the finest level of texture recovery. While for the invisible parts, the global features and the point-level features, where more coherent 3D-aware information are encoded, should contribute more. To model such complex feature relations, we employ the self-attention module  $\text{Attn}(\cdot)$  [70, 18, 20] with three attention heads for effective feature fusion, *i.e.*,

$$\text{Attn}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V, \quad (4)$$

where query  $Q$ , key  $K$  and value  $V$  are obtained by projecting the hierarchical features using MLPs.  $d$  is the scaling coefficient. The fused features are input into the NeRF decoder to predict the density  $\sigma$  and RGB  $\mathbf{c}$  for the point  $\mathbf{x}^t$ .

### 3.5. Training Details

SHERF contains five trainable modules, which are three hierarchical feature extraction modules, the feature fusion transformer and the NeRF decoder, which are trained in an

\*If the L2 norm distance is large than a threshold 0.05, we directly set the alpha  $\sigma$  and color  $\mathbf{c}$  of 3D point as a small value, *e.g.*,  $\sigma = -80$  and  $\mathbf{0}$ .



Table 1: Performance (PSNR, SSIM and LPIPS) comparison among NHP, MPS-NeRF and our SHERF method on the THuman, RenderPeople, ZJU\_MoCap and HuMMan datasets.

Method	THuman						RenderPeople					
	Novel View			Novel Pose			Novel View			Novel Pose		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
PixelNeRF [81]	16.51	0.65	0.35	-	-	-	-	-	-	-	-	-
NHP [39]	22.53	0.88	0.17	20.25	0.86	0.19	20.59	0.81	0.22	19.60	0.77	0.25
MPS-NeRF [23]	21.72	0.87	0.18	21.68	0.87	0.18	20.72	0.81	0.24	20.19	0.80	0.25
<b>SHERF (Ours)</b>	<b>24.66</b>	<b>0.91</b>	<b>0.10</b>	<b>24.26</b>	<b>0.91</b>	<b>0.11</b>	<b>22.88</b>	<b>0.88</b>	<b>0.14</b>	<b>21.98</b>	<b>0.86</b>	<b>0.15</b>

Method	ZJU_MoCap						HuMMan					
	Novel View			Novel Pose			Novel View			Novel Pose		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
NHP [39]	21.66	0.87	0.17	21.57	0.87	0.17	18.99	0.84	0.18	18.32	0.83	0.18
MPS-NeRF [23]	21.86	0.87	0.17	21.60	0.87	0.17	17.44	0.82	0.19	17.43	0.82	0.19
<b>SHERF (Ours)</b>	<b>22.87</b>	<b>0.89</b>	<b>0.12</b>	<b>22.38</b>	<b>0.89</b>	<b>0.12</b>	<b>20.83</b>	<b>0.89</b>	<b>0.12</b>	<b>20.43</b>	<b>0.88</b>	<b>0.11</b>

end-to-end manner. During training, for the same actor, we randomly sample image pairs from target and input views. By inputting the input view images to the above described process, we aim to reconstruct the actor in the target view. Four loss functions are used to supervise the training.

**Photometric Loss.** Given the ground truth target image  $C(\mathbf{r})$  and predicted image  $\hat{C}(\mathbf{r})$ , we apply the photometric loss as follows:

$$\mathcal{L}_{color} = \frac{1}{|\mathcal{R}|} \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{C}(\mathbf{r}) - C(\mathbf{r})\|_2^2, \quad (5)$$

where  $\mathcal{R}$  denotes the set of rays, and  $|\mathcal{R}|$  is the number of rays in  $\mathcal{R}$ .

**Mask Loss.** We also leverage the human region masks for Human NeRF optimization. The mask loss is defined as:

$$\mathcal{L}_{mask} = \frac{1}{|\mathcal{R}|} \sum_{\mathbf{r} \in \mathcal{R}} \|\hat{M}(\mathbf{r}) - M(\mathbf{r})\|_2^2, \quad (6)$$

where  $\hat{M}(\mathbf{r})$  is the accumulated volume density and  $M(\mathbf{r})$  is the ground truth binary mask label.

**SSIM Loss.** We further employ SSIM to ensure the structural similarity between ground truth and synthesized images, *i.e.*,

$$\mathcal{L}_{SSIM} = \text{SSIM}(\hat{C}(\mathbf{r}), C(\mathbf{r})). \quad (7)$$

**LPIPS Loss.** The perceptual loss LPIPS is also utilized to ensure the quality of rendered image, *i.e.*,

$$\mathcal{L}_{LPIPS} = \text{LPIPS}(\hat{C}(\mathbf{r}), C(\mathbf{r})). \quad (8)$$

In summary, the overall loss function contains four components, *i.e.*,

$$\mathcal{L} = \mathcal{L}_{color} + \lambda_1 \mathcal{L}_{mask} + \lambda_2 \mathcal{L}_{SSIM} + \lambda_3 \mathcal{L}_{LPIPS}, \quad (9)$$

where  $\lambda$ 's are the loss weights. Empirically, we select  $\lambda_1 = 0.1$ ,  $\lambda_2 = \lambda_3 = 0.01$  to ensure the same magnitude for each loss term.

## 4. Experiments

### 4.1. Experimental Setup

**Datasets.** Four large-scale human datasets are used for evaluation, *i.e.*, THuman [86], RenderPeople [1], ZJU\_MoCap [58] and HuMMan [9]. For ZJU\_MoCap, 9 subjects are split into 6 for training and 3 for testing. For each training and test subject, 100 frames are sampled for training or evaluation. For THuman, we randomly select 90 subjects as the training set and 10 subjects for testing. For each subject, we randomly sample 20 frames for training or evaluation. For RenderPeople, we randomly sample 450 subjects as the training set and 30 subjects for testing. For each subject, we randomly sample 10 frames for training or evaluation. For HuMMan, we use 317 sequences as the training set and 22 sequences for testing, following the official data split (HuMMan-Recon) for human reconstruction studies. For each sequence, we randomly sample 17 frames for training or evaluation.

**Comparison Methods.** To the best of our knowledge, we are the first to study the setting of single-image generalizable and animatable Human NeRF. We adapt two state-of-the-art generalizable Human NeRF methods designed for multi-view settings, *i.e.*, NHP [39] and MPS-NeRF [23] to our setting. For fair evaluation, we also compare with PixelNeRF [81], which is a generalizable NeRF reconstruction methods using single images as inputs. PixelNeRF is not specifically designed for human. Therefore, we only evaluate it for novel view synthesis.

**Implementation Details.** Our 2D Encoder adopts a pre-trained ResNet18 backbone to extract 1D vector  $\mathbf{f} \in \mathcal{R}^{512}$  for global feature extraction, and feature maps  $\mathbf{f} \in \mathcal{R}^{64 \times 256 \times 256}$  for point-level and pixel-aligned feature extraction. To preserve more low-level information, we further perform positional encoding to the RGB values and



Figure 3: Novel view and novel pose synthesis results produced by NHP, MPS-NeRF and SHERF on RenderPeople, THuman and ZJU\_MoCap. Zoom in for the best view.

append them to 2D feature maps to form feature maps  $f \in \mathcal{R}^{96 \times 256 \times 256}$ . The Mapping Network and Style-Based Encoder are adopted from EG3D [11]. Four layers of sparse convolutions [16] are used to extract 96-dimensional point-level features. The queried global feature, point-level feature and pixel-aligned feature are concatenated and projected to 32 channels before fed into the feature fusion transformer. The feature fusion transformer contains one self-attention layer with three heads. NeRF decoder is the same as that used in [23]. For LBS and inverse LBS of SMPL, we use the transformation or inverse transformation matrix of the nearest SMPL vertex.

**Evaluation Metrics.** To quantitatively evaluate the quality of rendered novel view and novel pose images, we report the peak signal-to-noise ratio (PSNR) [66], structural similarity index (SSIM) [74] and Learned Perceptual Image Patch Similarity (LPIPS) [83]. Instead of computing the metrics for

the whole image, we follow previous Human NeRF methods [59, 23] to project the 3D human bounding box to each camera plane to obtain the bounding box mask and report these metrics based on the masked area.

## 4.2. Quantitative Results

As shown in Tab. 1, SHERF significantly outperforms NHP and MPS-NeRF in all evaluation metrics on all four datasets. NHP and MPS-NeRF, as the SOTA generalizable Human NeRF methods in multi-view human image input setting, achieve reasonable performance in the novel view synthesis task. However, both NHP and MPS-NeRF focus on the local feature extraction and lack the ability to complement information missing from partial inputs, explaining their poor performance. They also fail to have good performance in the novel pose synthesis task, especially for NHP, which models neural radiance field in the target space.

Table 2: Ablation study on THuman. The left side shows different design components that are ablated on.

Global Feature	Point-Level Feature	Pixel-Aligned Feature	Feature Fusion	Novel View			Novel Pose		
				PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
✓				22.38	0.89	0.14	22.35	0.89	0.14
✓	✓			23.26	0.89	0.13	23.03	0.89	0.14
✓		✓		23.72	0.90	0.12	23.65	0.90	0.12
	✓	✓		24.08	0.90	0.12	23.73	0.90	0.12
✓	✓	✓		24.44	<b>0.91</b>	0.11	24.08	<b>0.91</b>	<b>0.11</b>
✓	✓	✓	✓	<b>24.66</b>	<b>0.91</b>	<b>0.10</b>	<b>24.26</b>	<b>0.91</b>	<b>0.11</b>

In contrast, SHERF achieves the best performance on both novel view synthesis and novel pose synthesis tasks. In addition to these three metrics, we also perform a user study and report human’s preference scores on rendered images. As shown in Fig. 4, SHERF has a clear advantage over two baseline methods.

### 4.3. Qualitative Results

We show the rendering images of our SHERF and two baseline methods in Fig. 3. NHP and MPS-NeRF produce reasonable RGB renderings in novel views, but they fail to recover details, *e.g.*, face details and cloth patterns. Thanks to the bank of hierarchical features, our SHERF successfully recovers face details and cloth patterns by enhancing the information from the input 2D observation and complementing the information missing from the input image. For example, SHERF renders the 3D human cloth with the same patterns as the input image and the 3D human back cloth with reasonable colors which are not observable from the input image. In novel pose synthesis, NHP produces distorted rendering results as it models the neural radiance filed in the target space. Compared with MPS-NeRF, SHERF shows better rendering results in novel pose synthesis.

### 4.4. Ablation Study

To validate the effectiveness of the proposed hierarchical feature extraction components and feature fusion transformer, we subsequently add different components and evaluate the performance on the THuman dataset. The results are reported on Tab. 2 and one visualization example is shown in Fig. 5. Given a single image input, only using global features can render images with reasonable RGB colors on test subjects. Adding point-level features or pixel-aligned features can further improve the image quality. Qualitatively, pixel-aligned features preserves more fine-grained details, *e.g.*, the cloth wrinkle in Fig. 5. Combining both point-Level and pixel-aligned features with global features can further improve the performance and render images with correct colors and fine-grained details, erasing the small artifacts when only point-level and pixel-aligned features are used




Figure 4: User preference scores on rendering results.




Figure 5: Qualitative results of ablation study on THuman. Refer to red arrows to see the cloth color and wrinkle difference, and purple arrows to see the erasion of black artifacts.

(see purple arrows in Fig. 5). Finally, using the feature fusion transformer, we further improve the reconstruction quality.

### 4.5. Further Analysis

**Discussion on Training Protocols.** Previous generalizable Human NeRF methods [23] carefully pick input camera views and fix them during training. Switching to a more challenging setting of single image input, a straightforward way is to extend this training setting by only using the front view as input. We evaluate the model trained with this setting with different settings. As shown in Tab. 3, when evaluated in either setting, SHERF leads the performance. Especially, model shows high rendering result when evaluated with the front view input setting. However, in most real-world scenarios, human images are captured from random camera viewing angles, challenging the above front view input training setting. Therefore, we further evaluate the above model using random view inputs. As shown in the right half of Tab. 3, the performance is much worse than front view inputs. It is also worse than the model trained with free view inputs in Tab. 1. The conclusion is that instead of a fixed viewing angle input, our free view input training setting is a more reasonable choice for real-world scenarios.

**Analysis on Different Viewing Angles as Inputs.** Further evaluations are performed to better understand how SHERF would perform given images with different viewing angles. We evaluate SHERF and baseline methods with input images from 12 camera viewing angles, which are evenly sampled from  $[0^\circ, 360^\circ]$ . For each input view, the target is to render the other 11 views, on which the mean PSNR is calculated

Table 3: Performance (PSNR) comparison among NHP, MPS-NeRF and SHERF trained with front view input and evaluated with different settings on THuman dataset.

Method	Front View Input		Free View Input	
	Novel View	Novel Pose	Novel View	Novel Pose
NHP [39]	24.00	19.75	20.59	19.17
MPS-NeRF [23]	23.29	23.15	21.56	21.46
<b>SHERF (Ours)</b>	<b>24.63</b>	<b>24.05</b>	<b>22.60</b>	<b>22.36</b>

and reported in Fig. 6(a). We find that SHERF 1) consistently outperforms SOTA baseline methods on all viewing angle inputs, and 2) shows robust performance to different viewing angle inputs.

**Analysis on Viewing Angle Difference Between Target and Observation.** As shown in Fig. 6(b), we study the effect of viewing angle difference between targets and inputs on the novel view synthesis. Two main trends can be found. 1) The smaller the viewing angle difference is, the easier for models to perform novel view synthesis. 2) Across all input settings, SHERF consistently outperforms baseline methods. By refining the indicators, we provide more detailed evaluation and gain more insight into the models.




Figure 6: **Analysis on Input Viewing Angles.** (a) reports novel view synthesis PSNR with different viewing angles as inputs. (b) shows novel view synthesis PSNR with different viewing angle difference between targets and inputs.

**Generalizability Analysis.** To further study the generalization ability, we directly inference NHP, MPS-NeRF and SHERF models, which are pre-trained on THuman, on ZJU\_Mocap test sets. As shown in Tab. 4, without additional fine-tuning, SHERF achieves better performance than NHP and MPS-NeRF. SHERF even shows comparable performance with NHP and MPS-NeRF that are trained on ZJU\_MoCap. To further show the generalizability on in-the-wild images, we evaluate SHERF on DeepFashion [45] images with SMPL and camera estimated from a 2D image using CLIFF [41] and show novel view and pose synthesis visualization results in Fig. 7.

**Comparison with PIFu.** For PIFu series, as PIFuHD can-

Table 4: Generalization ability comparison by cross-validating NHP, MPS-NeRF and SHERF trained with THuman on ZJU-MoCap.

Method	Novel View			Novel Pose		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
NHP [39]	<b>22.07</b>	0.88	0.16	20.70	0.86	0.18
MPS-NeRF [23]	20.36	0.85	0.17	19.96	0.85	0.17
<b>SHERF (Ours)</b>	21.87	<b>0.89</b>	<b>0.11</b>	<b>21.50</b>	<b>0.88</b>	<b>0.12</b>



Figure 7: Generalization ability of SHERF on in-the-wild DeepFashion images.

not predict 3D human with textures, we opt to compare with PIFu, which predicts textured mesh from single images. We evaluate PIFu on RenderPeople using the official models and compare with SHERF in terms of visualization results. Thanks to the volume rendering in NeRF and direct supervision on images, SHERF can generate textures with higher fidelity and better details than PIFu as shown in Fig. 8. Moreover, 3D human reconstructed by PIFu cannot be animated, while SHERF can naturally generalize to novel poses.



Figure 8: Visualization results comparison between PIFu and our SHERF.

**Comparison with modulation and adaptive normalization.** Even though global, point-level and pixel-aligned features contain different levels of semantics, they share the same shape, and is better suited for a transformer fusion module, where the learnable attention map allows sufficient

flexibility. We compare our transformer fusion module with other two classic feature fusion methods, *i.e.*, AdaIN and SIREN. We report their PSNR in novel view/pose synthesis in Tab. 5, which shows our transformer fusion is the most effective one.

Table 5: Performance (PSNR, SSIM and LPIPS) comparison among modulation, adaptive normalization and our Transformer fusion on the RenderPeople dataset.

Method	Novel View			Novel Pose		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
AdaIN	22.38	0.87	0.16	21.65	0.85	0.17
SIREN	21.62	0.85	0.18	21.09	0.84	0.19
<b>Transformer (Ours)</b>	<b>22.88</b>	<b>0.88</b>	<b>0.14</b>	<b>21.98</b>	<b>0.86</b>	<b>0.15</b>

**Runtime Analysis.** One big advantage of generalizable Human NeRF is that the reconstruction happens in only one forward pass. The inference time is essential for down-stream applications. Therefore, we also report and compare the inference frames per second (FPS) of SHERF and baseline methods in Tab. 6.

Table 6: Runtime comparison between NHP, MPS-NeRF and SHERF. ‘‘FPS’’ represents inference frames per second, the higher the better.

Method	NHP	MPS-NeRF	SHERF
FPS	0.15	0.60	<b>1.33</b>

## 5. Discussion and Conclusion

To conclude, we propose **SHERF**, the first generalizable human NeRF model that recovers animatable 3D humans from single human image inputs. To render high-fidelity 3D humans, SHERF proposes to learn both global and local details from the bank of 3D-aware hierarchical features comprising global features, point-level features, and pixel-aligned features. By using a feature fusion transformer, SHERF successfully enhances the information from the 2D observation and complements the information missing from the input image. On four large-scale human datasets, SHERF achieves state-of-the-art performance and renders high-fidelity images in both novel views and poses.

**Limitations:** 1) There still exists visible artifacts in target renderings when some body parts are occluded in the observation space. A better feature presentation like occlusion-aware features may be explored to solve this issue. 2) How to complement the information missing from single image input remains a challenging problem. SHERF starts from the reconstruction view and can only render deterministic

results when predicting novel views. One potential direction is to investigate the use of conditional generative models to diversely generate higher quality novel views.

**Potential Negative Societal Impacts:** SHERF can be misused to create fake images or videos of real humans and cause negative social impacts.

## 6. Acknowledgment

This study is supported by the Ministry of Education, Singapore, under its MOE AcRF Tier 2 (MOE-T2EP20221-0012), NTU NAP, and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).



## References

- [1] Renderpeople. In <https://renderpeople.com/3d-people>, 2018.
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *2018 International Conference on 3D Vision (3DV)*, pages 98–109. IEEE, 2018.
- [3] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8387–8397, 2018.
- [4] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2293–2303, 2019.
- [5] Thiemo 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.
- [6] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Combining implicit function learning and parametric models for 3d human reconstruction. In *European Conference on Computer Vision*, pages 311–329. Springer, 2020.
- [7] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Loopreg: Self-supervised learning of implicit surface correspondences, pose and shape for 3d human mesh registration. *Advances in Neural Information Processing Systems*, 33:12909–12922, 2020.
- [8] Aljaz Bozic, Pablo Palafox, Michael Zollhofer, Justus Thies, Angela Dai, and Matthias Nießner. Neural deformation graphs for globally-consistent non-rigid reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1450–1459, 2021.
- [9] 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 *17th European Conference on Computer Vision, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, pages 557–577. Springer, 2022.
- [10] Ang Cao, Chris Rockwell, and Justin Johnson. Fwd: Real-time novel view synthesis with forward warping and depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15713–15724, 2022.
- [11] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *arXiv*, 2021.
- [12] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021.
- [13] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, Xu Jia, and Huchuan Lu. Animatable neural radiance fields from monocular rgb videos. *arXiv preprint arXiv:2106.13629*, 2021.
- [14] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [15] Hongsuk Choi, Gyeongsik Moon, Matthieu Armando, Vincent Leroy, Kyoung Mu Lee, and Gregory Rogez. Mononhr: Monocular neural human renderer. *arXiv preprint arXiv:2210.00627*, 2022.
- [16] Sponv Contributors. Sponv: Spatially sparse convolution library. <https://github.com/traveller59/sponv>, 2022.
- [17] Enric Corona, Mihai Zanfir, Thiemo Alldieck, Eduard Gabriel Bazavan, Andrei Zanfir, and Cristian Sminchisescu. Structured 3d features for reconstructing relightable and animatable avatars. *arXiv preprint arXiv:2212.06820*, 2022.
- [18] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [19] Zijian Dong, Chen Guo, Jie Song, Xu Chen, Andreas Geiger, and Otmar Hilliges. Pina: Learning a personalized implicit neural avatar from a single rgb-d video sequence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20470–20480, 2022.
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [21] Emilien Dupont, Bautista Miguel Angel, Alex Colburn, Aditya Sankar, Carlos Guestrin, Josh Susskind, and Qi Shan. Equivariant neural rendering. *arXiv preprint arXiv:2006.07630*, 2020.
- [22] Chen Gao, Yichang Shih, Wei-Sheng Lai, Chia-Kai Liang, and Jia-Bin Huang. Portrait neural radiance fields from a single image. *arXiv preprint arXiv:2012.05903*, 2020.
- [23] Xiangjun Gao, Jiaolong Yang, Jongyoo Kim, Sida Peng, Zicheng Liu, and Xin Tong. Mps-nerf: Generalizable 3d human rendering from multiview images. *arXiv preprint arXiv:2203.16875*, 2022.
- [24] Pengsheng Guo, Miguel Angel Bautista, Alex Colburn, Liang Yang, Daniel Ulbricht, Joshua M Susskind, and Qi Shan. Fast and explicit neural view synthesis. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3791–3800, 2022.
- [25] Tong He, John Collomosse, Hailin Jin, and Stefano Soatto. Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. *Advances in Neural Information Processing Systems*, 33:9276–9287, 2020.
- [26] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. Arch++: Animation-ready clothed human recon-

- struction revisited. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11046–11056, 2021.
- [27] 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. *arXiv*, 2022.
- [28] 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.
- [29] Wonbong Jang and Lourdes Agapito. Codenerf: Disentangled neural radiance fields for object categories. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12949–12958, 2021.
- [30] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. Selfrecon: Self reconstruction your digital avatar from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5605–5615, 2022.
- [31] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. *arXiv preprint arXiv:2203.12575*, 2022.
- [32] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8320–8329, 2018.
- [33] James T Kajiya and Brian P Von Herzen. Ray tracing volume densities. *ACM SIGGRAPH computer graphics*, 18(3):165–174, 1984.
- [34] 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.
- [35] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020.
- [36] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5253–5263, 2020.
- [37] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. Pare: Part attention regressor for 3d human body estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11127–11137, 2021.
- [38] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2252–2261, 2019.
- [39] 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.
- [40] Ruilong Li, Yuliang Xiu, Shunsuke Saito, Zeng Huang, Kyle Olszewski, and Hao Li. Monocular real-time volumetric performance capture. In *European Conference on Computer Vision*, pages 49–67. Springer, 2020.
- [41] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. Cliff: Carrying location information in full frames into human pose and shape estimation. In *European Conference on Computer Vision*, pages 590–606. Springer, 2022.
- [42] Zhe Li, Tao Yu, Chuanyu Pan, Zerong Zheng, and Yebin Liu. Robust 3d self-portraits in seconds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1344–1353, 2020.
- [43] Kai-En Lin, Lin Yen-Chen, Wei-Sheng Lai, Tsung-Yi Lin, Yi-Chang Shih, and Ravi Ramamoorthi. Vision transformer for nerf-based view synthesis from a single input image. *arXiv preprint arXiv:2207.05736*, 2022.
- [44] Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou, and Wenping Wang. Neural rays for occlusion-aware image-based rendering. In *CVPR*, 2022.
- [45] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1096–1104, 2016.
- [46] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015.
- [47] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995.
- [48] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [49] Lu Mi, Abhijit Kundu, David Ross, Frank Dellaert, Noah Snavely, and Alireza Fathi. im2nerf: Image to neural radiance field in the wild. *arXiv preprint arXiv:2209.04061*, 2022.
- [50] 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. In *European conference on computer vision*, pages 405–421. Springer, 2020.
- [51] 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. In *European conference on computer vision*, pages 405–421. Springer, 2020.
- [52] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022.
- [53] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In

- Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [54] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5762–5772, 2021.
- [55] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [56] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10975–10985, 2019.
- [57] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Animatable neural radiance fields for human body modeling. *arXiv e-prints*, pages arXiv–2105, 2021.
- [58] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021.
- [59] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021.
- [60] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J Black. Clothcap: Seamless 4d clothing capture and retargeting. *ACM Transactions on Graphics (ToG)*, 36(4):1–15, 2017.
- [61] Konstantinos Rematas, Ricardo Martin-Brualla, and Vittorio Ferrari. Sharf: Shape-conditioned radiance fields from a single view. *arXiv preprint arXiv:2102.08860*, 2021.
- [62] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *arXiv preprint arXiv:2201.02610*, 2022.
- [63] 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.
- [64] 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.
- [65] Mehdi SM Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lučić, Daniel Duckworth, Alexey Dosovitskiy, et al. Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6229–6238, 2022.
- [66] Umme Sara, Morium Akter, and Mohammad Shorif Uddin. Image quality assessment through fsim, ssim, mse and psnr—a comparative study. *Journal of Computer and Communications*, 7(3):8–18, 2019.
- [67] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems*, 32, 2019.
- [68] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. *Advances in Neural Information Processing Systems*, 34:12278–12291, 2021.
- [69] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Yifan Wang, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. *arXiv preprint arXiv:2111.05849*, 2021.
- [70] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [71] Dan Wang, Xinrui Cui, Septimiu Salcudean, and Z Jane Wang. Generalizable neural radiance fields for novel view synthesis with transformer. *arXiv preprint arXiv:2206.05375*, 2022.
- [72] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snively, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021.
- [73] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdf. In *European conference on computer vision*, volume 4, 2022.
- [74] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [75] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humanerf: Free-viewpoint rendering of moving people from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16210–16220, 2022.
- [76] 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.
- [77] 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.
- [78] Hongyi Xu, Thiemo Alldieck, and Cristian Sminchisescu. H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. *Advances in Neural Information Processing Systems*, 34, 2021.

- [79] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Ghum & ghuml: Generative 3d human shape and articulated pose models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6184–6193, 2020.
- [80] Ze Yang, Shenlong Wang, Sivabalan Manivasagam, Zeng Huang, Wei-Chiu Ma, Xinchun Yan, Ersin Yumer, and Raquel Urtasun. S3: Neural shape, skeleton, and skinning fields for 3d human modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13284–13293, 2021.
- [81] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021.
- [82] 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR2021)*, June 2021.
- [83] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [84] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. Humannerf: Generalizable neural human radiance field from sparse inputs. *arXiv preprint arXiv:2112.02789*, 2021.
- [85] 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.
- [86] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7739–7749, 2019.
- [87] Hao Zhu, Xinxin Zuo, Sen Wang, Xun Cao, and Ruigang Yang. Detailed human shape estimation from a single image by hierarchical mesh deformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4491–4500, 2019.