

# MonoHuman: Animatable Human Neural Field from Monocular Video


Zhengming Yu<sup>1</sup>, Wei Cheng<sup>1,2</sup>, Xian Liu<sup>3</sup>, Wayne Wu<sup>2</sup>, Kwan-Yee Lin<sup>2,3</sup> ✉

<sup>1</sup>SenseTime Research <sup>2</sup>Shanghai AI Laboratory

<sup>3</sup>The Chinese University of Hong Kong

yuhuaijin36@gmail.com chengwei@sensetime.com

alvinliu@ie.cuhk.edu.hk wuwenyan0503@gmail.com junyilin@cuhk.edu.hk



**Input: Monocular Video**

**Output: Animatable Avatar**

Figure 1. **Overview of MonoHuman.** We present MonoHuman, which learns an animatable human neural field from monocular videos. With monocular video input, our method can create an animatable human avatar that supports novel view synthesis and novel pose animation. Our code, models and demos are available at <https://yzmblog.github.io/projects/MonoHuman>.

## Abstract

Animating virtual avatars with free-view control is crucial for various applications like virtual reality and digital entertainment. Previous studies have attempted to utilize the representation power of the neural radiance field (NeRF) to reconstruct the human body from monocular videos. Recent works propose to graft a deformation network into the NeRF to further model the dynamics of the human neural field for animating vivid human motions. However, such pipelines either rely on pose-dependent representations or fall short of motion coherency due to frame-independent optimization, making it difficult to generalize to unseen pose sequences realistically. In this paper, we propose a novel framework **MonoHuman**, which robustly renders view-consistent and high-fidelity avatars under arbitrary novel poses. Our key insight is to model the deformation field with bi-directional constraints and explicitly leverage the off-the-peg keyframe information to reason the feature correlations for coherent results. Specifically, we first propose a Shared Bidirectional Deformation module, which creates a pose-independent generalizable deformation field by disentangling backward and forward deformation correspondences into shared skeletal motion weight and separate non-rigid motions. Then, we devise a Forward Correspondence Search module, which queries the correspondence feature of keyframes to guide the rendering net-

work. The rendered results are thus multi-view consistent with high fidelity, even under challenging novel pose settings. Extensive experiments demonstrate the superiority of our proposed MonoHuman over state-of-the-art methods.

## 1. Introduction

Rendering a free-viewpoint photo-realistic view synthesis of a digital avatar with explicit pose control is an important task that will bring benefits to AR/VR applications, virtual try-on, movie production, telepresence, etc. However, previous methods [32, 33, 58] usually require carefully-collected multi-view videos with complicated systems and controlled studios, which limits the usage in general and personalized scenarios applications. Therefore, though challenging, it has a significant application value to directly recover and animate the digital avatar from a monocular video.

Previous rendering methods [33] can synthesize realistic novel view images of the human body, but hard to animate the avatar in unseen poses. To address this, some recent methods deform the neural radiance fields (NeRF [27]) [32, 49] to learn a backward skinning weight of parametric models depending on the pose or individual frame index. They can animate the recovered human in novel poses with small variations to the training set. However, as the

blending weights are pose-dependent, these methods usually over-fit to the seen poses of training data, therefore, lack of generalizability [49]. Notably, one category of approaches [5, 21] solves this problem by learning a pose-independent forward blend weight in canonical space and using the root-finding algorithm to search the backward correspondence. However, the root-finding algorithm is time-consuming. [4] proposes to add a forward mapping network to help the learning of backward mapping, though with consistent constraint, their backward warping weights still are frame dependent. Some other works [16, 32, 42] leverage the blend weight from the template model like SMPL [25]. The accuracy of their deformation heavily hinges on the template model and always fails in the cloth parts as SMPL does not model it. How to learn an accurate and generalizable deformation field is still an open problem.

To reconstruct and animate a photo-realistic avatar that can generalize to unseen poses from the monocular video, we conclude three key observations to achieve generalizability from recent studies: **1)** The deformation weight field should be defined in canonical space and as pose-independent as possible [5, 48, 49]; **2)** An ideal deformation field should unify forward and backward deformation to alleviate ambiguous correspondence on novel poses; **3)** Direct appearance reference from input observation helps improve the fidelity of rendering.

We propose **MonoHuman**, a novel framework that enjoys the above strengths to reconstruct an animatable digital avatar from only monocular video. Concretely, we show how to learn more correct and general deformation from such limited video data and how to render realistic results from deformed points in canonical space. We first introduce a novel Shared Bidirectional Deformation module to graft into the neural radiance fields, which disentangles the backward and forward deformation into one shared skeletal motion and two separate residual non-rigid motions. The shared motion basis encourages the network to learn more general rigid transformation weights. The separate residual non-rigid motions guarantee the expressiveness of the accurate recovery of pose-dependent deformation. With the accurate learned deformation, we further build an observation bank consisting of information from sparse keyframes, and present a Forward Correspondence Search module to search observation correspondence from the bank which helps create high fidelity and natural human appearance, especially the invisible part in the current frame. With the above designs, our framework can synthesize a human at any viewpoint and any pose with natural shape and appearance.

To summarize, our main contributions are three-fold:

- We present a new approach MonoHuman that can synthesize the free viewpoint and novel pose sequences of a performer with explicit pose control, only requiring a monocular video as supervision.

- We propose the Shared Bidirectional Deformation module to achieve generalizable consistent forward and backward deformation, and the Forward Search Correspondence module to query correspondence appearance features to guide the rendering step.
- Extensive experiments demonstrate that our framework MonoHuman renders high-fidelity results and outperforms state-of-the-art methods.

## 2. Related Work

**Human Performance Capture.** Previous works reconstruct the human body geometry from multi-view videos [9, 40, 51] or depth cameras [28, 39, 53, 56] and the albedo map of surface mesh [11, 12]. Recent works model human geometry as radiance fields [6, 16, 19, 30–33, 49, 58] or distance functions [45, 52]. NeuralBody [32] uses structured pose features generated from SMPL [25] to condition the radiance field, enabling it to recover human performers and produce free-viewpoint images from sparse multi-view videos. Researchers further improve the generalization [6, 19, 58] ability or reconstruction quality [52] following the same setup. Although they have improved the rendering quality well, the application scenario of these methods is limited to mostly indoor multi-camera setups with accurate camera calibration and restricted capture volume.

**Human Rendering from Monocular Video.** To remove the multi-view constraints, many methods explore reconstructing the digital human from a single image or monocular video. [36, 37, 50] recover static clothed 3d human from a single image by learning the accurate surface reconstruction. In order to model the human dynamics, Some works [13, 54] learn to deform the pre-scanned human model from monocular video. [1] reconstructs the full human body from a self-rotating monocular video by optimizing the displacement of the SMPL model. Recently, SelfRecon [15] improves the reconstruction quality by representing the body motion as a learnable non-rigid motion along with the pre-defined SMPL skinning. Other works [30, 31, 35] model the dynamic human via deformable NeRF with a deformation or blending network that blends the ray from current motion space to a canonical space. However, the depth ambiguity of monocular video and insufficient pose observation may lead the deformation field to overfit to training data which further causes floating artifacts in novel views. To address these problems, some works [16, 48, 49] introduce motion priors to regularize the deformation. NeuMan [16] learns the blending field by extending the deformation weight of the closest correspondence from the SMPL mesh and optimizing it during training. HumanNeRF [49] decouples the deformation field as a skeleton motion field and a non-rigid motion field. These methods target at rendering free-viewpoint human images

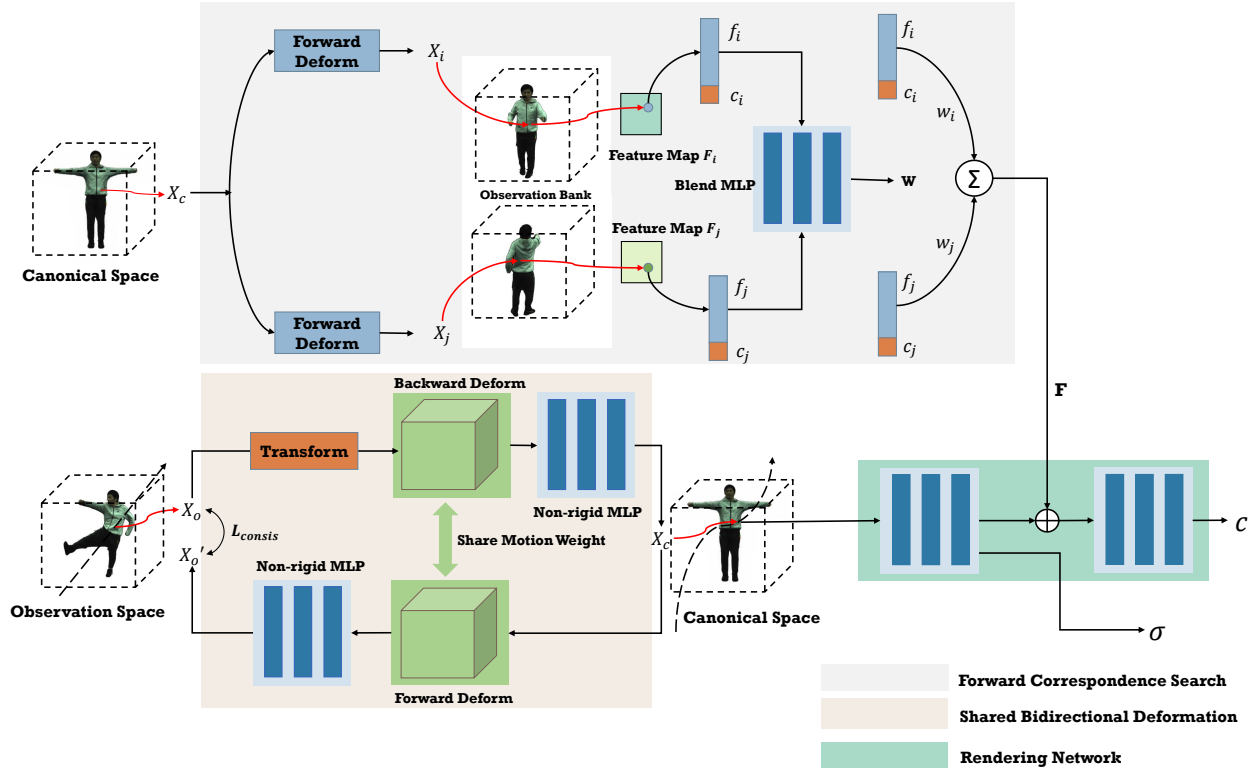


Figure 2. **Framework of MonoHuman.** In the Shared Bidirectional Deformation module, points in different observation space  $X_o$  will be deformed to  $X_c$  in the same canonical space by backward deforming  $D_b$ . And  $X_c$  can be deformed back into the points  $X'_o$  in observation space by symmetrical forward deformation  $D_f$ . A consistent loss  $L_{cons}$  is used to regularize the motion weight. In Forward Correspondence Search module,  $X_c$  is deformed to  $X_i$  and  $X_j$  in two correspondence space to get the feature  $f_i, f_j$  and the color  $c_i, c_j$ . These features are blended to the feature  $F$  by the weights Blend MLP generated.  $F$  is used to guide the rendering network finally.

for a bullet-time effect, while ours focuses on animating reconstructed avatars to out-of-distribution poses even challenge poses generated by text prompts.

**Human Animation.** Neural Actor [23] proposes deformable NeRF from UV space and refines the detail with the predicted texture map. [32] create an animatable model by conditioning the NeRF with pose-dependent latent code. [29, 41] drive the canonical model by transforming points to local bone coordinates. Moco-Flow [4] uses a complementary time-conditioned forward deformation network with cycle consistency loss to regularize the deformation field. These methods usually lean by memorizing the observed poses and are hard to generalize to novel poses. Recent works [5, 38] address this problem by combining a forward deformation network with the root-finding algorithm to learn the deformation field. However, the root-finding algorithm they use is highly time-consuming.

**Image-based rendering.** To synthesize novel view images without recovering detailed 3D geometry, previous works [8, 10] get novel view images from interpolating the light field. Though achieving realistic results, the view range they can synthesize is limited. To address this problem, some works [2, 34] seek the help of proxy geometry by inferring the depth map. Although they extend the render-

able view range, their generated results are sensitive to the accuracy of the proxy geometry. With the development of deep learning techniques, many works [7, 14, 17, 20, 24, 44, 46] introduce learnable components to the image-based rendering methods and improve the robustness. Specifically, IBNet [46] learns the blending weight of image features from a sparse set of nearby views. NeuRay [24] further predicts the visibility of 3D points to input views within their representation. These works show the powerful effect of image features from different views for rendering.

### 3. Our Approach

We present **MonoHuman** that reconstruct an animatable digital avatar from monocular video. The generated avatar can be controlled in free view-point and arbitrary novel poses. The pipeline is illustrated in Fig. 2. Specifically, we first formulate the problem in Sec. 3.1, and then the *Shared Bidirectional Deformation Module* which deforms points between observation space and canonical space is elaborated in Sec. 3.2. Sec. 3.3 introduces the *Forward Correspondence Search Module* to extract the correspondence features in keyframes. Finally, the training objectives and volume rendering process are described in Sec. 3.4.

### 3.1. Preliminaries and Problem Setting

Given an image sequence of a person and corresponding poses, segmentation masks, as well as the camera’s intrinsic and extrinsic parameters, following [32, 49] we use a neural field to represent the human from observation space in one canonical space. Specifically, the color and density of point  $x_o$  in observation space are model as:

$$(c(\mathbf{x}_o), \sigma(\mathbf{x}_o)) = F_c(D(\mathbf{x}_o, \mathbf{p})), \quad (1)$$

where  $D$  is a backward deformation mapping that takes the body pose  $\mathbf{p}$  and point in observation space  $\mathbf{x}_o$  as input, outputs the point in the canonical space  $\mathbf{x}_c$ .  $F_c$  is a mapping network that takes point in the canonical space  $\mathbf{x}_c$  as input and outputs its color value  $c$  and density  $\sigma$ . With this representation, the final rendering step using volume rendering [26, 27] is done in one canonical space.

### 3.2. Shared Bidirectional Deformation Module

To avoid the overfitting problem mentioned in Sec. 1, we define the deformation field in canonical space following HumanNeRF [49]. However, the proposed single backward deformation is only constrained by the image reconstruction loss, which is under-constrained and will gain more deform errors when it meets unseen poses. We show this in phenomenon Sec. 4. Intuitively, a loss that is defined only in the deformation field will help to regularize the deformation field. But how to build such a constraint is non-trivial. MoCo-Flow [4] uses an additional time-conditioned forward deformation MLP to constrain the consistency of deformation. BANMo [55] uses different MLPs to generate pose-conditioned deformation weights to achieve forward and backward deformation. However, as the deformation fields are two different MLPs and frame-dependent or pose-dependent, they still suffer from the over-fitting problem. Inspired by the works above, we design our Shared Bidirectional Deformation module with a forward and backward deform using the same motion weight defined in canonical space. Since our design is pose-independent, it helps the model’s generalization to out-of-distribution poses. In practice, we formulate the backward deformation as:

$$D_b : (\mathbf{x}_o, \mathbf{p}) \rightarrow \mathbf{x}_c, \quad (2)$$

where  $\mathbf{p} = \{\omega_i\}$  are  $K$  joints rotations represented as axis-angle vectors. Similar to [49], the complete deformation consists of two parts: motion weight deformation and non-rigid deformation:

$$D_b(\mathbf{x}_o, \mathbf{p}) = T_{mo}^b(\mathbf{x}_o, \mathbf{p}) + T_{NR}^b(T_{mo}^b(\mathbf{x}_o, \mathbf{p}), \mathbf{p}), \quad (3)$$

where computation of  $T_{mo}^b$  is similar to linear blend skinning:

$$T_{mo}^b(\mathbf{x}_o, \mathbf{p}) = \sum_{i=1}^K w_o^i(\mathbf{x}_o)(R_i \mathbf{x}_o + t_i), \quad (4)$$

in which  $w_o^i$  is the blend weight of  $i$ -th bone, and  $R_i, t_i$  are the rotation and translation that map the bone’s coordinates from observation to canonical space. We compute  $w_o^i(\mathbf{x}_o)$  by a set of motion weight volumes defined in canonical space  $w_c^i$ :

$$w_o^i(\mathbf{x}_o) = \frac{w_c^i(R_i \mathbf{x}_o + t_i)}{\sum_{k=1}^K w_c^k(R_k \mathbf{x}_o + t_k)}. \quad (5)$$

We solve the volume  $W_c(x_c) = \{w_c^i(x_c)\}$  using CNN which generates from a random constant latent code  $\mathbf{z}$  same as HumanNeRF [49]:

$$W_c(\mathbf{x}_c) = \text{CNN}(\mathbf{x}_c; \mathbf{z}). \quad (6)$$

We additionally use the motion weight defined in canonical space to achieve forward deformation:

$$D_f : (\mathbf{x}_c, \mathbf{p}) \rightarrow \mathbf{x}_o, \quad (7)$$

different from backward deformation, the forward motion weight can be directly queried by  $\mathbf{x}_o$  as:

$$T_{mo}^f(\mathbf{x}_c, \mathbf{p}) = \sum_{i=1}^K w_c^i(\mathbf{x}_c) \mathbf{x}_c, \quad (8)$$

because the motion weight is defined in canonical space. As for non-rigid deformation, we use another MLP to compute forward deform. Both forward and backward pose-dependent deformation can be written in the same formula:

$$T_{NR}(\mathbf{x}, \mathbf{p}) = \text{MLP}_{NR}(T_{mo}(\mathbf{x}, \mathbf{p}), \mathbf{p}), \quad (9)$$

In order to add the constraint that only related to the deformation field as a regularization, we use the intuition of consistency of forward and backward deformation, the consistent loss  $\mathbf{L}_{consis}$  is computed as:

$$L_{consis} = \begin{cases} d & \text{if } d \geq \theta \\ 0 & \text{else} \end{cases} \quad d = L_2(\mathbf{x}_o, D_f(D_b(\mathbf{x}_o, \mathbf{p}), \mathbf{p})), \quad (10)$$

where  $L_2$  means the  $L_2$  distance calculation, and it only penalize the points whose  $L_2$  distance is greater than threshold  $\theta$  we set to avoid over regularization.

### 3.3. Forward Correspondence Search Module

Previous multi-view works [3, 46, 58] utilize the correspondence features from observed views at the same time stamp to guide the novel view synthesis. As our input images are from monocular video, we don’t have multi-view images at the same time. However, we can utilize the forward deformation in Sec. 3.2 to achieve dynamic deformation and find correspondences in different time stamps in our monocular video. Inspired by the previous works, We design our observation bank consisting of correspondence



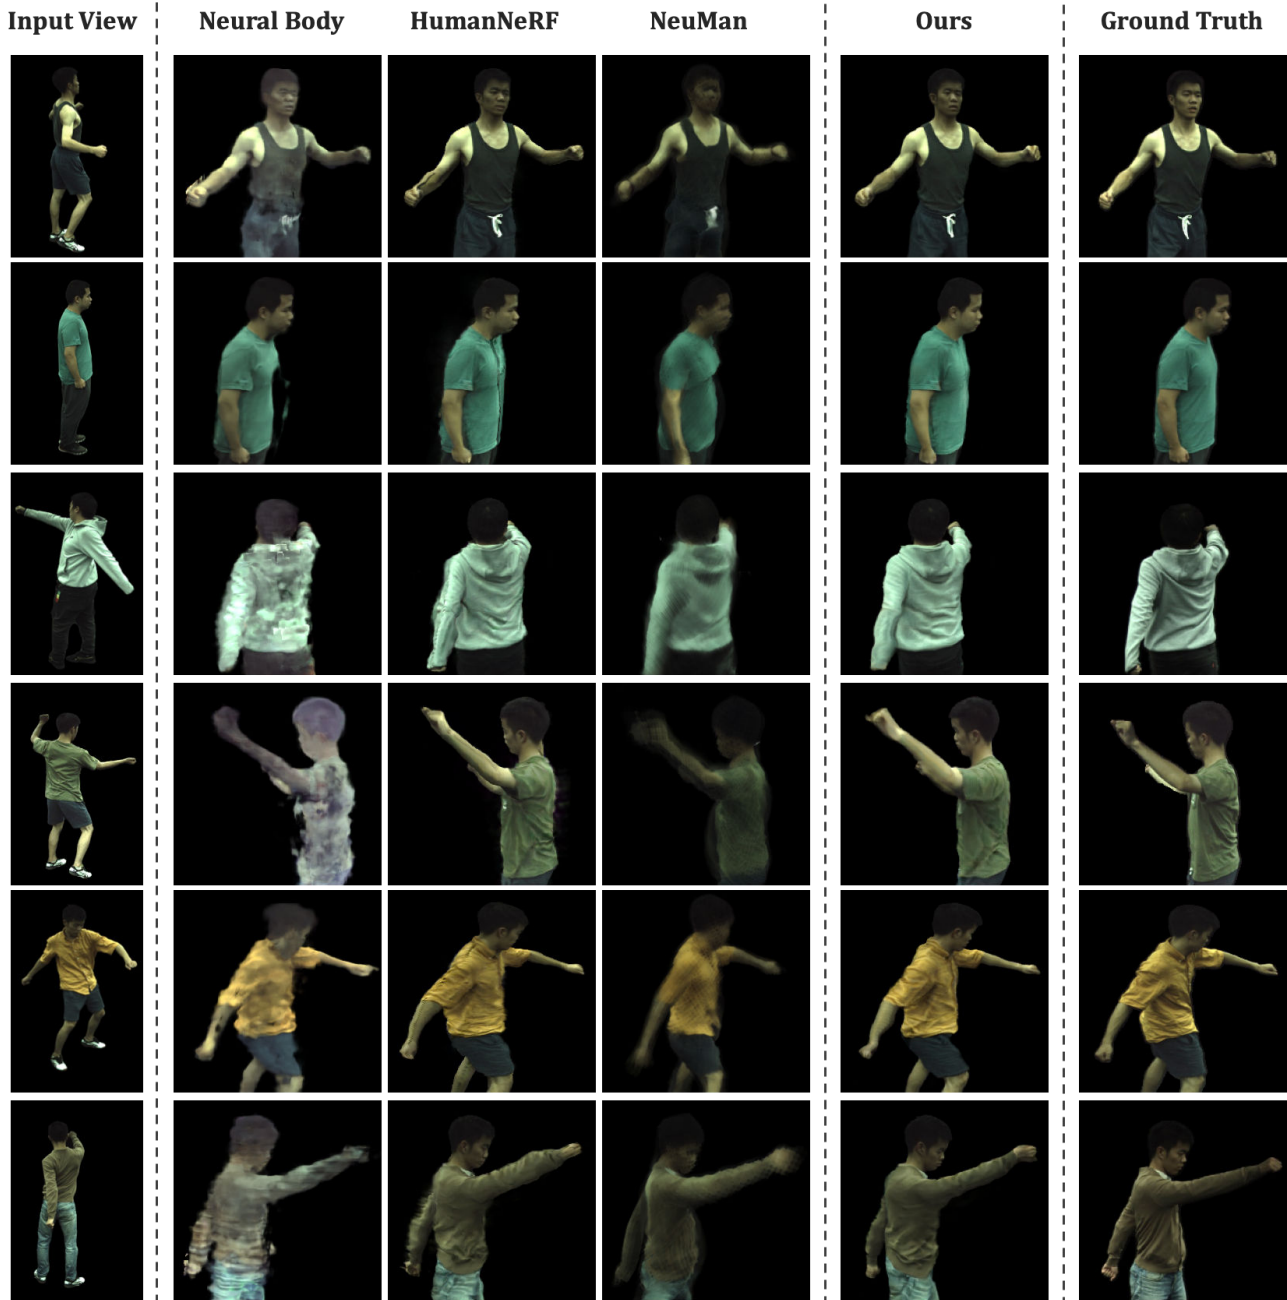


Figure 3. **Qualitative result of novel view setting in ZJU-MoCap.** We compare the novel view synthesis quality with baseline methods in ZJU-Mocap. Result shows that our method synthesizes more realistic images in novel view.

features to guide the rendering. We propose to build an observation bank by spanning the time and searching the correspondence features in these posed keyframe images from the input monocular video sequence in order to guide the novel view and novel pose synthesis. As illustrated in Fig. 4, we first separate the frames into two subsets containing the front and back view images respectively, according to the orientation of the pelvis. Then, we find the  $k$  pairs with the closest pose from these two sets. We reconstruct the texture map for the  $k$  pairs, and finally select the pair

with the highest texture map complementarity. However, as the lack of accurate camera parameters, searching the correspondence in the monocular input setting is a non-trivial problem. To address this, we utilize the advantage of forward deformation described in Sec. 3.2 to search the correspondence features of point  $x_c$  in canonical space. Specifically, we first deform  $x_c$  to  $x_o^i$  from canonical to correspondence space using Eq. 7:

$$\mathbf{x}_o^i = D_f(\mathbf{x}_c, \mathbf{p}_i), \quad (11)$$

where  $\mathbf{p}_i$  denotes the pose in  $i$ -th frame. Then with the camera parameters of  $i$ -th frame, we project the point  $x_o^i$  to the image coordinate:

$$\mathbf{x}_i = K_i E_i \mathbf{x}_o^i, \quad (12)$$

where  $K_i, E_i$  are the intrinsic and extrinsic camera parameters of  $i$ -th frame respectively.  $x_i$  is the pixel location in  $i$ -th frame, which is used to sample the feature  $f_i$  and color  $c_i$  of  $i$ -th image. We follow the U-Net feature extract from IBNet [46] to extract the image features. We get the feature of  $j$ -th frame  $f_j$  and  $c_j$  in the same way. Then we use a Blend MLP to map these two features to the blend weights:

$$\begin{aligned} \mathbf{F} &= w_i(f_i; c_i) + w_j(f_j; c_j), \\ w &= \text{MLP}_{\text{blend}}((f_i; c_i), (f_j; c_j)), \end{aligned} \quad (13)$$

where  $(;)$  is the concatenation of two vectors. The blended feature  $\mathbf{F}$  is used to guide the rendering network.




Figure 4. **The Illustration of Keyframe Selection.** We select the keyframes according to their pose similarity and texture map complementarity.




Figure 5. **Qualitative results on novel view from Internet video.** Top: invisible (video source), Down: Cha-Cha (video source)

### 3.4. Volume Rendering and Network Training

**Volume rendering with deformation.** We conduct the rendering step in one canonical space. With the guidance of the feature extracted by the Forward Correspondence Search

module described in Sec. 3.3, our rendering network takes input canonical space point  $\mathbf{x}_i$  and feature  $\mathbf{F}$ , and outputs its color  $\mathbf{c}(\mathbf{x}_i)$  and density  $\sigma(\mathbf{x}_i)$ . We denote our whole rendering network as a mapping  $F_r$ :

$$\mathbf{c}(\mathbf{x}_i), \sigma(\mathbf{x}_i) = F_r(\gamma(\mathbf{x}_i), \mathbf{F}), \quad (14)$$

where  $\gamma$  is a sinusoidal positional encoding [27] of  $x_i$ . We finally render the neural field in canonical space using the volume rendering equation [26,27]. The color  $C(\mathbf{r})$  of a ray  $\mathbf{r}$  with  $D$  samples can be written as:

$$\begin{aligned} C(\mathbf{r}) &= \sum_{i=1}^D \left( \prod_{j=1}^{i-1} (1 - \alpha_j) \right) \alpha_i \mathbf{c}(\mathbf{x}_i), \\ \alpha_i &= 1 - \exp(-\sigma(\mathbf{x}_i) \Delta t_i), \end{aligned} \quad (15)$$

where  $\Delta t_i$  is the interval between sample  $i$  and  $i + 1$ .

**Network training.** We optimize our network with  $\mathcal{L}_{\text{MSE}}$ ,  $\mathcal{L}_{\text{LPIPS}}$  and  $\mathcal{L}_{\text{CONSIS}}$ :

$$\mathcal{L}_{\text{MSE}} = \sum_{\mathbf{r} \in R} \left( \|C(\mathbf{r}) - \hat{C}(\mathbf{r})\|_2^2 \right) \quad (16)$$

where  $C(\mathbf{r}), \hat{C}(\mathbf{r})$  are rendered color and ground truth color of camera ray  $\mathbf{r}$ . We also adopt a perceptual loss LPIPS [57] in our final loss. The consistent loss described in Sec. 3.2 is used for optimization, with the final loss function as:

$$\mathcal{L} = \mathcal{L}_{\text{MSE}} + \lambda \mathcal{L}_{\text{LPIPS}} + \mathcal{L}_{\text{CONSIS}} \quad (17)$$

where  $\lambda$  is a weight balance coefficient.

## 4. Experiments

### 4.1. Dataset and Preprocessing

We use ZJU-MoCap dataset [33] and in-the-wild video collected from the Internet to evaluate our method. We follow [49] to select the same 6 subjects for our evaluation. As ZJU-MoCap is carefully collected in an indoor environment, we directly use the annotations in the dataset. For the wild video we gathered from the Internet, we used PARE [18] to estimate the camera matrix and body pose, and RVM [22] to get the segmentation mask. We select keyframes as illustrated in Fig. 4 and Sec. 3.3.

### 4.2. Experimental Settings

**Comparison baselines.** We compare our method with (1) NeuralBody [33], which uses structured latent code to represent the human body; (2) HumanNeRF [49], which achieves state-of-the-art image synthesis performance of digital human by learning a motion field mapping network from monocular video. (3) NeuMan [16], who reconstructs the scene and human using two different NeRFs, we only compare with NeuMan’s human NeRF without background.

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS* $\downarrow$
Neural Body [33]	28.72	0.9611	52.25
HumanNeRF [49]	29.61	0.9625	38.45
NeuMan [16]	28.96	0.9479	60.74
<b>Ours</b>	<b>30.26</b>	<b>0.9692</b>	<b>30.92</b>

Table 1. **Novel view synthesis quantitative comparison on ZJU-MoCap dataset.** We compute averages over 6 sequences. For results of each subject please refer to supplementary material. Note that in all table results, LPIPS\* = LPIPS  $\times 10^3$ .

**Implementation details.** We adopt a U-Net to extract image features from keyframes similar to [46], and the deform error thread  $\theta = 0.05$ . During training, the learning rate of the Shared Bidirectional Deformation module is set to  $5 \times 10^{-6}$  and  $5 \times 10^{-5}$  for the rest. The overall framework trains on a single V100 GPU for 70 hours.

### 4.3. Quantitative Evaluation

**Evaluation metrics.** We use PSNR and SSIM [47] to evaluate the generated image quality. As PSNR prefers smooth results but may with bad visual quality [49, 57], we additionally adopt LPIPS [57] to measure the perceived distance between the synthesis image and ground truth image.

**Comparison settings.** For comparison of the performance of novel view synthesis and animation, we follow [49] to only use the camera 1 data in ZJU-MoCap for training. We further split the frames in camera 1 as seen-pose Set A and unseen-pose Set B in 4:1 ratio. We compare our method with baselines in two settings. 1) The novel view setting in Tab. 1, where we follow [49] to use frames in Set A as training and the synchronous frames in the other 22 cameras as evaluation. 2) The novel pose setting in Tab. 2, where we use frames in Set A as training, the frames in Set B, and its synchronous frames in other 22 cameras as evaluation.

**Evaluation results.** The results of the novel view setting and novel pose setting are shown in Tab. 1 and Tab. 2, respectively. We can see that our MonoHuman framework outperforms existing methods in most metrics in both settings. With the help of correspondence features and the consistent constraint of the deformation field, we can synthesize more realistic results. Though the PSNR metric of NeuralBody is not bad, they synthesize poor visual quality images in both novel view and novel pose (as PSNR prefers smooth results [57], but it can be verified by LPIPS). Our improvement in PSNR over HumanNeRF [49] is not significant, and slightly lower in subject 387 when testing in the novel pose. But in LIPIS, our method has about 19.6% and 9.46% improvement in novel view and novel pose setting respectively. Note that novel pose synthesis is more challenging than novel view synthesis, but in some subjects, the evaluation metrics behave the opposite because some poses are repeated in the data. So we conduct extract qualitative comparison in out of distribution pose. NeuMan [16] only

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS* $\downarrow$
Neural Body [33]	28.54	0.9604	53.91
HumanNeRF [49]	29.74	0.9655	34.79
NeuMan [16]	28.75	0.9406	62.35
<b>Ours</b>	<b>30.05</b>	<b>0.9684</b>	<b>31.51</b>

Table 2. **Novel pose synthesis quantitative comparison on ZJU-MoCap dataset.** We compute averages over 6 sequences. For results of each subject please refer to supplementary material.

updates the SMPL parameter during training, and models the deformation as the closest point’s LBS weight, which leads to poor performance in novel views.

### 4.4. Qualitative Evaluation

To compare the generated results, we visualize the novel view synthesis results in ZJU-MoCap dataset of NeuralBody, HumanNeRF and our method in Fig.3. NeuralBody tends to generate vague images with large noise in novel view. Both HumanNeRF and our method synthesize realistic images in novel view but predict artifacts in some detail areas. We also show the novel view synthesis results in videos collected from the Internet in Fig.5. For extreme views, HumanNeRF tends to generate artifacts in clothes and face areas. while our methods can retain these detail due to the correct deform and help with guided features.

As the motions in ZJU-MoCap and self-collected videos are simple and usually repetitive, we need to find a way to evaluate the animation ability in some complex motions. To do this, we use the MDM [43], which can generate complex motions corresponding to the text input. We use such a model to generate challenge motions like back-flip and martial pose sequences to animate the avatar reconstructed by HumanNeRF and our method. As shown in Fig. 6, HumanNeRF predicts multiple artifacts in some extreme poses like squatting and jumping high. It also fails to deform the arms correctly in punch motions and back-flip poses in the air. With the help of the Shared Bidirectional Deformation Module, we proposed, we can handle the deformation in such challenging poses and generate more realistic results.

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS* $\downarrow$
Neural Body [33]	28.72	0.9611	52.25
HumanNeRF [49]	29.61	0.9625	38.45
NeuMan [16]	28.96	0.9479	60.74
Ours (w/o $\mathcal{L}_{\text{CONSIS}}$ )	30.17	0.9689	31.76
Ours (w/o feat)	30.16	0.9689	31.39
<b>Ours (full model)</b>	<b>30.26</b>	<b>0.9692</b>	<b>30.92</b>

Table 3. **Ablation study on ZJU-MoCap in novel view setting.** We compute averages over 6 sequences. LPIPS\* = LPIPS  $\times 10^3$ .



Text input: The person is doing backflip




Figure 6. **Qualitative results on challenge poses generated by MDM [43] through text input.** We evaluate our method driven by challenge pose sequence generated by MDM model.

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS* $\downarrow$
Neural Body [33]	28.54	0.9604	53.91
HumanNeRF [49]	29.74	0.9655	34.79
NeuMan [16]	28.75	0.9406	62.35
Ours (w/o $\mathcal{L}_{\text{CONSIS}}$ )	29.86	0.9679	32.39
Ours (w/o feat)	29.88	0.9681	31.80
<b>Ours (full model)</b>	<b>30.05</b>	<b>0.9684</b>	<b>31.51</b>

Table 4. **Ablation study on ZJU-MoCap in novel pose setting.** We compute averages over 6 sequences. LPIPS\* = LPIPS  $\times 10^3$ .

#### 4.5. Ablation Study

We conduct ablation studies for the two key modules in our model, the Shared Bidirectional Deformation module and the Forward Correspondence Search module in ZJU-MoCap dataset. To further explore the function of these two modules, we experiment in novel view and novel pose settings separately. The result is shown in Tab. 3 and Tab. 4 respectively. As shown in Tab. 3, these two modules contribute roughly the same to the result. Because the input poses in the novel view setting are all seen in training data, so even without the consistency constraints, the deformation is mostly correct. But when it comes to the novel poses, LPIPS is higher as the deformation is more difficult in the novel view setting. Furthermore, the incorrect deformation will influence the correspondence features we searched by our FCS module, which makes the generated result worse.

The influence without the Forward Correspondence Search module is nearly the same in both novel view and novel pose settings, as the guidance of correspondence fea-

ture is equally important in both novel view and novel pose settings. However, the performance of our FCS module relies on the Shared Bidirectional Deformation module, and the result of Tab. 4 verifies this property. In comparison to other baselines, these two modules we proposed are helpful to the generated results. We furthermore visualize the effectiveness of our consistent loss, please refer to the supplementary material for more details.

## 5. Discussion

We propose MonoHuman, which robustly renders view-consistent avatars at novel poses of high fidelity. We propose a novel Shared Bidirectional Deformation module to handle out-of-distribution novel pose deformation. Furthermore, a Forward Correspondence Search module which queries the correspondence feature to guide the rendering network is used to help generate photo-realistic results. To evaluate our approach, we generate challenging poses by Human Motion Diffusion Model [43] to animate our avatar. The result shows that our model can generate high-fidelity images even under challenging novel pose settings.

**Limitations and Future work.** Our methods can create an animatable avatar with only a monocular video. However, our method still has a few limitations: 1) Our synthesis result relies on the accuracy of the pose and mask annotations. It would be interesting to explore how to make pose and mask correction during training. 2) Our method is trained in a case-specific manner. Studying how to generalize different people with monocular videos will be valuable.



## References

- [1] 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. 2
- [2] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM Transactions on Graphics (TOG)*, 2013. 3
- [3] 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. 4
- [4] Xuelin Chen, Weiyu Li, Daniel Cohen-Or, Niloy J Mitra, and Baoquan Chen. Moco-flow: Neural motion consensus flow for dynamic humans in stationary monocular cameras. In *Computer Graphics Forum*, volume 41, pages 147–161. Wiley Online Library, 2022. 2, 3, 4
- [5] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. SNARF: Differentiable forward skinning for animating non-rigid neural implicit shapes. *ICCV*, 2021. 2, 3
- [6] Wei Cheng, Su Xu, Jingtian Piao, Chen Qian, Wayne Wu, Kwan-Yee Lin, and Hongsheng Li. Generalizable neural performer: Learning robust radiance fields for human novel view synthesis. *arXiv preprint arXiv:2204.11798*, 2022. 2
- [7] Inchang Choi, Orazio Gallo, Alejandro Troccoli, Min H Kim, and Jan Kautz. Extreme view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7781–7790, 2019. 3
- [8] Abe Davis, Marc Levoy, and Fredo Durand. Unstructured light fields. In *Computer Graphics Forum*, volume 31, pages 305–314. Wiley Online Library, 2012. 3
- [9] Edilson De Aguiar, Carsten Stoll, Christian Theobalt, Naveed Ahmed, Hans-Peter Seidel, and Sebastian Thrun. Performance capture from sparse multi-view video. In *ACM SIGGRAPH 2008 papers*, pages 1–10. 2008. 2
- [10] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. *SIGGRAPH*, 1996. 3
- [11] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, et al. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics (TOG)*, 2019. 2
- [12] Kaiwen Guo, Feng Xu, Tao Yu, Xiaoyang Liu, Qionghai Dai, and Yebin Liu. Real-time geometry, albedo, and motion reconstruction using a single rgb-d camera. *ACM Transactions on Graphics (ToG)*, 36(4):1, 2017. 2
- [13] Marc Habermann, Weipeng Xu, Michael Zollhofer, Gerard Pons-Moll, and Christian Theobalt. Deepcap: Monocular human performance capture using weak supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5052–5063, 2020. 2
- [14] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (TOG)*, 37(6):1–15, 2018. 3
- [15] 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. 2
- [16] 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. 2, 6, 7, 8
- [17] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi. Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 35(6):1–10, 2016. 3
- [18] 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. 6
- [19] 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
- [20] Youngjoong Kwon, Stefano Petrangeli, Dahun Kim, Hao-liang Wang, Eunbyung Park, Viswanathan Swaminathan, and Henry Fuchs. Rotationally-temporally consistent novel view synthesis of human performance video. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 387–402. Springer, 2020. 3
- [21] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhofer, Jurgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. *arXiv preprint arXiv:2206.08929*, 2022. 2
- [22] Shanchuan Lin, Linjie Yang, Imran Saleemi, and Soumyadip Sengupta. Robust high-resolution video matting with temporal guidance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 238–247, 2022. 6
- [23] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Trans. Graph.(ACM SIGGRAPH Asia)*, 2021. 3
- [24] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7824–7833, 2022. 3
- [25] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 2015. 2

- [26] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1995. 4, 6
- [27] 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. *ECCV*, 2020. 1, 4, 6
- [28] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 343–352, 2015. 2
- [29] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. *ICCV*, 2021. 3
- [30] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *ICCV*, 2021. 2
- [31] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. HyperNeRF: A higher-dimensional representation for topologically varying neural radiance fields. *SIGGRAPH Asia*, 2021. 2
- [32] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. *ICCV*, 2021. 1, 2, 3, 4
- [33] 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. *CVPR*, 2021. 1, 2, 6, 7, 8
- [34] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)*, 36(6):1–11, 2017. 3
- [35] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. *CVPR*, 2020. 2
- [36] 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. 2
- [37] 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. 2
- [38] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2886–2897, 2021. 3
- [39] Ari Shapiro, Andrew Feng, Ruizhe Wang, Hao Li, Mark Bolas, Gerard Medioni, and Evan Suma. Rapid avatar capture and simulation using commodity depth sensors. *Computer Animation and Virtual Worlds*, 25(3-4):201–211, 2014. 2
- [40] Leonid Sigal, Alexandru O Balan, and Michael J Black. Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International journal of computer vision*, 87(1):4–27, 2010. 2
- [41] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. In *Advances in Neural Information Processing Systems*, 2021. 3
- [42] Gusi Te, Xiu Li, Xiao Li, Jinglu Wang, Wei Hu, and Yan Lu. Neural capture of animatable 3d human from monocular video. *arXiv preprint arXiv:2208.08728*, 2022. 2
- [43] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 7, 8
- [44] Justus Thies, Michael Zollhöfer, Christian Theobalt, Marc Stamminger, and Matthias Nießner. Ignor: Image-guided neural object rendering. *arXiv preprint arXiv:1811.10720*, 2018. 3
- [45] Garvita Tiwari, Dimitrije Antić, Jan Eric Lenssen, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Pose-ndf: Modeling human pose manifolds with neural distance fields. In *European Conference on Computer Vision*, pages 572–589. Springer, 2022. 2
- [46] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. IBRNet: Learning multi-view image-based rendering. *CVPR*, 2021. 3, 4, 6, 7
- [47] 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. 7
- [48] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Vid2Actor: Free-viewpoint animatable person synthesis from video in the wild. *arXiv:2012.12884*, 2020. 2
- [49] 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. 1, 2, 4, 6, 7, 8
- [50] 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
- [51] Feng Xu, Yebin Liu, Carsten Stoll, James Tompkin, Gaurav Bharaj, Qionghai Dai, Hans-Peter Seidel, Jan Kautz, and Christian Theobalt. Video-based characters: creating new human performances from a multi-view video database. *SIGGRAPH*, 2011. 2
- [52] 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. 2

- [53] Lan Xu, Wei Cheng, Kaiwen Guo, Lei Han, Yebin Liu, and Lu Fang. Flyfusion: Realtime dynamic scene reconstruction using a flying depth camera. *IEEE transactions on visualization and computer graphics*, 27(1):68–82, 2019. [2](#)
- [54] Weipeng Xu, Avishek Chatterjee, Michael Zollhöfer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian Theobalt. Monoperfcap: Human performance capture from monocular video. *ACM Transactions on Graphics (ToG)*, 37(2):1–15, 2018. [2](#)
- [55] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2863–2873, June 2022. [4](#)
- [56] Tao Yu, Zerong Zheng, Kaiwen Guo, Jianhui Zhao, Qionghai Dai, Hao Li, Gerard Pons-Moll, and Yebin Liu. Doublefusion: Real-time capture of human performances with inner body shapes from a single depth sensor. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7287–7296, 2018. [2](#)
- [57] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *CVPR*, 2018. [6](#), [7](#)
- [58] 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. [1](#), [2](#), [4](#)