

# Animatable Neural Radiance Fields for Modeling Dynamic Human Bodies

Sida Peng<sup>1\*</sup> Junting Dong<sup>1\*</sup> Qianqian Wang<sup>2</sup> Shangzhan Zhang<sup>1</sup>  
 Qing Shuai<sup>1</sup> Xiaowei Zhou<sup>1</sup> Hujun Bao<sup>1†</sup>  
<sup>1</sup>Zhejiang University <sup>2</sup>Cornell University

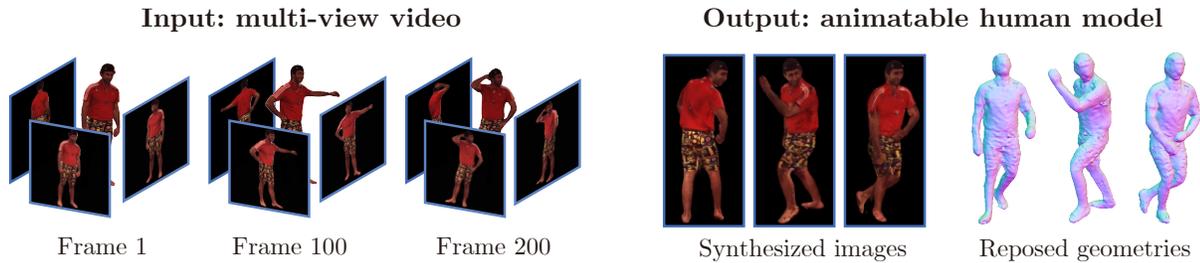


Figure 1: Given a multi-view video of a performer, our method reconstructs an animatable human model, which can be used for novel view synthesis and 3D shape generation under novel human poses.

## Abstract

*This paper addresses the challenge of reconstructing an animatable human model from a multi-view video. Some recent works have proposed to decompose a non-rigidly deforming scene into a canonical neural radiance field and a set of deformation fields that map observation-space points to the canonical space, thereby enabling them to learn the dynamic scene from images. However, they represent the deformation field as translational vector field or  $SE(3)$  field, which makes the optimization highly under-constrained. Moreover, these representations cannot be explicitly controlled by input motions. Instead, we introduce neural blend weight fields to produce the deformation fields. Based on the skeleton-driven deformation, blend weight fields are used with 3D human skeletons to generate observation-to-canonical and canonical-to-observation correspondences. Since 3D human skeletons are more observable, they can regularize the learning of deformation fields. Moreover, the learned blend weight fields can be combined with input skeletal motions to generate new deformation fields to animate the human model. Experiments show that our approach significantly outperforms recent human synthesis methods. The code and supplementary materials are available at [https://zju3dv.github.io/animatable\\_nerf/](https://zju3dv.github.io/animatable_nerf/).*

\*The first two authors contributed equally. The authors from Zhejiang University are affiliated with the State Key Lab of CAD&CG.

†Corresponding author: Hujun Bao.

## 1. Introduction

Rendering animatable human characters has a variety of applications such as free-viewpoint videos, telepresence, video games and movies. The core step is to reconstruct animatable human models, which tends to be expensive and time-consuming in traditional pipelines due to two factors. First, high-quality human reconstruction generally relies on complicated hardware, such as a dense array of cameras [56, 16] or depth sensors [10, 14]. Second, human animation requires skilled artists to manually create a skeleton suitable for the human model and carefully design skinning weights [29] to achieve realistic animation, which takes countless human labor.

In this work, we aim to reduce the cost of human reconstruction and animation, to enable the creation of digital humans at scale. Specifically, we focus on the problem of automatically reconstructing animatable humans from multi-view videos, as illustrated in Figure 1. However, this problem is extremely challenging. There are two core questions we need to answer: how to represent animatable human models and how to learn this representation from videos?

Recently, neural radiance fields (NeRF) [41] has proposed a representation that can be efficiently learned from images with a differentiable renderer. It represents static 3D scenes as color and density fields, which work particularly well with volume rendering techniques. To extend NeRF to handle non-rigidly deforming scenes, [46, 51] decompose a video into a canonical NeRF and a set of deformation fields that transform observation-space points at each video

frame to the canonical space. The deformation field is represented as translational vector field [51] or SE(3) field [46]. Although they can handle some dynamic scenes, they are not suited for representing animatable human models due to two reasons. First, jointly optimizing NeRF with translational vector fields or SE(3) fields without motion prior is an extremely under-constrained problem [51, 30]. Second, they cannot explicitly synthesize novel scenes given input motions for animation.

To overcome these problems, we propose a novel motion representation named neural blend weight field. Based on the skeleton-driven deformation framework [29], blend weight fields are combined with 3D human skeletons to generate deformation fields. This representation has two advantages. First, since the human skeleton is easy to track [22], it does not need to be jointly optimized and thus provides an effective regularization on the learning of deformation fields. Second, by learning an additional neural blend weight field at the canonical space, we can explicitly animate the neural radiance field with input motions.

We evaluate our approach on the H36M [19] and ZJU-MoCap [49] datasets that capture dynamic humans in complex motions with synchronized cameras. Across all video sequences, our approach exhibits state-of-the-art performances on novel view synthesis and novel pose synthesis. In addition, our method is able to reconstruct the 3D human shape at the canonical space and repose the geometry.

In summary, this work has the following contributions:

- We introduce a novel representation called neural blend weight field, which can be combined with NeRF and 3D human skeletons to recover animatable human models from multi-view videos.
- Our approach demonstrates significant performance improvement on novel view synthesis and novel pose synthesis compared to recent human synthesis methods on the H36M and ZJU-MoCap datasets.

## 2. Related work

**Human reconstruction.** Modeling human characters is the first step of traditional animation pipelines. To achieve high-quality reconstruction, most methods rely on complicated hardware [10, 14, 59, 11, 16]. Recently, some works [58, 44, 41, 32] have attempted to learn 3D representations from images with differentiable renderers, which reduces the number of input camera views and achieves impressive reconstruction results. However, they have difficulty in recovering reasonable 3D human shapes when the camera views are too sparse, as shown in [49]. Instead of optimizing the network parameters per scene, [42, 54, 67, 55] utilize networks to learn human shape priors from ground-truth 3D data, allowing them to reconstruct human shapes from even a single image.

**Human animation.** Skeletal animation [29, 25] is a common approach to animate human models. It first creates a scale-appropriate skeleton for the human mesh and then assigns each mesh vertex a blend weight that describes how the vertex position deforms with the skeleton. Skinned multi-person linear model (SMPL) [36] learns a skeleton regressor and blend weights from a large amount of ground-truth 3D meshes. Based on SMPL, some works [48, 24, 27, 21, 13] reconstruct an animated human mesh from sparse camera views. However, SMPL only describes the naked human body and thus cannot be directly used to render photorealistic images. To overcome this problem, [3, 2, 4] apply vertex displacements to the SMPL model to capture the human clothing and hair. [61] proposes a 2D warping method to deform the SMPL model to fit the input image. Recent implicit function-based methods [45, 40, 9] have exhibited state-of-the-art reconstruction quality. [18, 5] combine implicit function learning with the SMPL model to obtain detailed animatable human models. [12] combines a set of local implicit functions with human skeletons to represent dynamic humans. [64] proposes to animate occupancy networks with a linear blend skinning algorithm. However, these methods all need the supervision of 3D ground-truth data.

**Neural rendering.** To reduce the requirement for the reconstruction quality, some methods [57, 60, 34, 62, 28] improve the rendering pipeline with neural networks. Based on the advances in image-to-image translation techniques [20], [38, 8, 39] train a network to map 2D skeleton images to target rendering results. Although these methods can synthesize photorealistic images under novel human poses, they have difficulty in rendering novel views. To improve the performance of novel view synthesis, [57, 60, 62, 1, 50, 65, 52] introduce 3D representations into the rendering pipeline. [60] establishes neural texture maps and uses UV maps to obtain feature maps in the image space, which is then interpreted into images with a neural renderer. [62, 1] reconstruct a point cloud from input images and learn a 3D feature for each point. Then, they project 3D features into a 2D feature map and employ a network to render images. However, 2D convolutional networks have difficulty in rendering inter-view consistent images, as shown in [58].

To solve this problem, [35, 44, 41, 31, 33] interpret features into colors in 3D space and then accumulate them into 2D images. [35] uses 3D convolutional networks to produce discretized RGB- $\alpha$  volumes. Neural radiance fields (NeRF) [41] proposes to represent 3D scenes with color and density fields, which works well with the volumetric rendering and gives state-of-the-art performances on novel view synthesis. [49] combines NeRF with the SMPL model, allowing it to handle dynamic humans and synthesize photorealistic novel views from very sparse camera views.

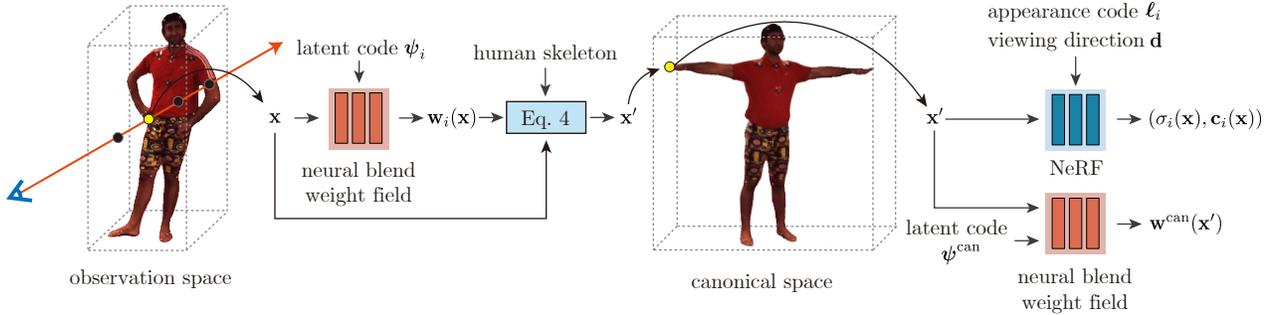


Figure 2: **Overview of our approach.** Given a query point  $\mathbf{x}$  in the observation space at frame  $i$ , we infer its blend weight  $\mathbf{w}_i(\mathbf{x})$  using a neural blend weight field that is conditioned on the latent code  $\psi_i$ . Based on the blend weight and the human skeleton, we can obtain the corresponding point  $\mathbf{x}'$  in the canonical space using equation (4). Taking the transformed point  $\mathbf{x}'$ , observation-space viewing direction  $\mathbf{d}$ , and appearance code  $\ell_i$  as inputs, the template NeRF model predicts the volume density and color. To animate the template NeRF, we also learn a neural blend field  $\mathbf{w}^{\text{can}}(\mathbf{x}')$  at the canonical space.

### 3. Method

Given a multi-view video of a performer, our task is to reconstruct an animatable human model that can be used to synthesize free-viewpoint videos of the performer under novel human poses. The cameras are synchronized and calibrated. For each frame, we assume the 3D human skeleton is given, which can be obtained with marker-based or marker-less pose estimation systems [19, 22]. For each image, [15] is used to extract the foreground human mask, and the values of the background image pixels are set to zero.

The overview of our approach is shown in Figure 2. We decompose a non-rigidly deforming human body into a canonical human model represented by a neural radiance field (Section 3.1) and a per-frame blend weight field (Section 3.2) that is used to establish correspondences between the observation space and canonical space. Then we discuss how to learn the representation on the multi-view video (Section 3.3). Based on blend weight fields, we are able to animate the canonical human model (Section 3.4).

#### 3.1. Representing videos with neural radiance fields

NeRF represents a static scene as a continuous volumetric representation. For any 3D point, it takes a spatial position  $\mathbf{x}$  and viewing direction  $\mathbf{d}$  as input to a neural network and outputs a volume density  $\sigma$  and color  $\mathbf{c}$ .

Inspired by [46, 51], we extend NeRF to represent the dynamic human body by introducing deformation fields, as shown in Figure 2. Specifically, for each video frame  $i \in \{1, \dots, N\}$ , we define a deformation field  $T_i$  that transforms observation-space points to the canonical space. Given the canonical-frame density model  $F_\sigma$ , the density model at frame  $i$  can be thus defined as:

$$(\sigma_i(\mathbf{x}), \mathbf{z}_i(\mathbf{x})) = F_\sigma(\gamma_{\mathbf{x}}(T_i(\mathbf{x}))), \quad (1)$$

where  $\mathbf{z}_i(\mathbf{x})$  is the shape feature in the original NeRF, and

$\gamma_{\mathbf{x}}$  is the positional encoding [41] for spatial location.

When predicting the color, we define a per-frame latent code  $\ell_i$  to encode the state of the human appearance in frame  $i$ . Similarly, with the canonical-frame color model  $F_c$ , the color model at frame  $i$  can be defined as:

$$\mathbf{c}_i(\mathbf{x}) = F_c(\mathbf{z}_i(\mathbf{x}), \gamma_{\mathbf{d}}(\mathbf{d}), \ell_i), \quad (2)$$

where  $\gamma_{\mathbf{d}}$  is the positional encoding for viewing direction.

There are several ways to represent the deformation field, such as translational vector field [51, 30] and SE(3) field [46]. However, as discussed in [46, 30], optimizing a radiance field together with a deformation field is an ill-posed problem that is prone to local optima. To overcome this problem, [46, 30] propose many regularization techniques to facilitate the training, which makes the optimization process complex. Moreover, their representations cannot robustly generate new deformation fields given novel motion sequences.

#### 3.2. Neural blend weight fields

Considering that we aim to model dynamic humans, it is natural to leverage the human priors to learn the deformation field, which helps us to solve the under-constrained problem. Specifically, we construct the deformation field based on the 3D human skeleton and the skeleton-driven deformation framework [29].

The human skeleton defines  $K$  parts, which produce  $K$  transformation matrices  $\{G_k\} \in SE(3)$ . The detailed derivation is listed in the supplementary material. In the linear blend skinning algorithm [29], a canonical-space point  $\mathbf{v}$  is transformed to the observation space using

$$\mathbf{v}' = \left( \sum_{k=1}^K w(\mathbf{v})_k G_k \right) \mathbf{v}, \quad (3)$$

where  $w(\mathbf{v})_k$  is the blend weight of  $k$ -th part. Similarly, for an observation-space point  $\mathbf{x}$ , if we know its corresponding

blend weights, we are able to transform it to the canonical space using

$$\mathbf{x}' = \left( \sum_{k=1}^K w^o(\mathbf{x})_k G_k \right)^{-1} \mathbf{x}, \quad (4)$$

where  $w^o(\mathbf{x})$  is the blend weight function defined in the observation space. To obtain the blend weight field, a natural idea is to define a function that maps a 3D point to blend weights, which then gives the dynamic radiance fields based on equations (1), (2) and (4). However, we find that jointly learning NeRF with the blend weight field is still ill-posed and is prone to local minima.

To solve this problem, we seek the human priors in 3D statistical body models [36, 53, 47, 63] to regularize the learned blend weights. Specifically, for any 3D point, we assign an initial blend weight based on the body model and then use a network to learn a residual vector, resulting in the neural blend weight field. In practice, the residual vector fields for all training video frames are implemented using a single MLP network  $F_{\Delta\mathbf{w}} : (\mathbf{x}, \psi_i) \rightarrow \Delta\mathbf{w}_i$ , where  $\psi_i$  is a per-frame learned latent code and  $\Delta\mathbf{w}_i$  is a vector  $\in \mathbb{R}^K$ . The neural blend weight field at frame  $i$  is defined as:

$$\mathbf{w}_i(\mathbf{x}) = \text{norm}(F_{\Delta\mathbf{w}}(\mathbf{x}, \psi_i) + \mathbf{w}^s(\mathbf{x}, S_i)), \quad (5)$$

where  $\mathbf{w}^s$  is the initial blend weights that are computed based on the statistical body model  $S_i$ , and we define  $\text{norm}(\mathbf{w}) = \mathbf{w} / \sum w_i$ . Without loss of generality, we adopt SMPL [36] as the body model, which can be obtained by fitting the SMPL model to the 3D human skeleton [22]. Note that this idea can also apply to other human models [53, 47, 63]. To compute  $\mathbf{w}^s$ , we take the strategy proposed in [18, 6]. For any 3D point, we first find the closest surface point on the SMPL mesh. Then, the target blend weight is computed by performing barycentric interpolation of the blend weights of three vertices on the corresponding mesh facet.

To animate the learned template NeRF, we additionally learn a neural blend weight field  $\mathbf{w}^{\text{can}}$  at the canonical space. The SMPL blend weight field  $\mathbf{w}^s$  is calculated using the T-pose SMPL model, and  $F_{\Delta\mathbf{w}}$  is conditioned on an additional latent code  $\psi^{\text{can}}$ . We utilize the inherent consistency between blend weights to optimize the neural blend weight field  $\mathbf{w}^{\text{can}}$ , which will be described in Section 3.3.

Instead of learning blend weight fields at both observation and canonical spaces, an alternative method is to only learn the blend weight field at the canonical space as in Equation (3), which specifies the canonical-to-observation correspondences. However, ‘‘inverting’’ Equation (3) to get observation-to-canonical correspondences for rendering is non-trivial. We would need to first build a dense set of observation-to-canonical correspondences by densely sampling points at the canonical space and evaluating their

blend weights. Then, for any observation-space point, we can interpolate its corresponding canonical point based on the pre-computed correspondences. This process is complex and time-consuming. Moreover, as the sampled points are discretized, the calculated correspondences tend to be coarse. In contrast, learning blend weights at observation spaces enables us to easily obtain the observation-to-canonical correspondences based on Equation (4).

### 3.3. Training

Based on the dynamic radiance field  $\sigma_i$  and  $\mathbf{c}_i$ , we can use volume rendering techniques [23, 41] to synthesize images of particular viewpoints for each video frame  $i$ . The near and far bounds of volume rendering are estimated by computing the 3D boxes that bound the SMPL meshes. The parameters of  $F_\sigma$ ,  $F_{\mathbf{c}}$ ,  $F_{\Delta\mathbf{w}}$ ,  $\{\ell_i\}$  and  $\{\psi_i\}$  are jointly optimized over the multi-view video by minimizing the difference between the rendered pixel color  $\tilde{\mathbf{C}}_i(r)$  and the observed pixel color  $\mathbf{C}_i(r)$ :

$$L_{\text{rgb}} = \sum_{r \in \mathcal{R}} \|\tilde{\mathbf{C}}_i(\mathbf{r}) - \mathbf{C}_i(\mathbf{r})\|_2, \quad (6)$$

where  $\mathcal{R}$  is the set of rays passing through image pixels.

To learn the neural blend weight field  $\mathbf{w}^{\text{can}}$  at the canonical space, we introduce a consistency loss between blend weight fields. As shown by equations (3) and (4), two corresponding points at canonical and observation spaces should have the same blend weights. For an observation-space point  $\mathbf{x}$  at frame  $i$ , we map it to the canonical-space point  $T_i(\mathbf{x})$  using equation (4). The consistency loss between blend weight fields is defined as:

$$L_{\text{nsf}} = \sum_{\mathbf{x} \in \mathcal{X}_i} \|\mathbf{w}_i(\mathbf{x}) - \mathbf{w}^{\text{can}}(T_i(\mathbf{x}))\|_1, \quad (7)$$

where  $\mathcal{X}_i$  is the set of 3D points sampled within the 3D human bounding box at frame  $i$ . The coefficient weights of  $L_{\text{rgb}}$  and  $L_{\text{nsf}}$  are both set to 1.

### 3.4. Animation

**Image synthesis.** To synthesize images of the performer under novel human poses, we similarly construct the deformation fields that transform the 3D points to the canonical space. Given a novel human pose, our method updates the pose parameters in the SMPL model and computes the SMPL blend weight field  $\mathbf{w}^s$  based on the new parameters  $S^{\text{new}}$ . Then, the neural blend weight field  $\mathbf{w}^{\text{new}}$  for the novel human pose is defined as:

$$\mathbf{w}^{\text{new}}(\mathbf{x}, \psi^{\text{new}}) = \text{norm}(F_{\Delta\mathbf{w}}(\mathbf{x}, \psi^{\text{new}}) + \mathbf{w}^s(\mathbf{x}, S^{\text{new}})), \quad (8)$$

where the  $F_{\Delta\mathbf{w}}$  is conditioned on a new latent code  $\psi^{\text{new}}$ . Based on the  $\mathbf{w}^{\text{new}}$  and equation (4), we can generate the

deformation field  $T^{\text{new}}$  for the novel human pose. The parameters of  $\psi^{\text{new}}$  are optimized using

$$L_{\text{new}} = \sum_{\mathbf{x} \in \mathcal{X}^{\text{new}}} \|\mathbf{w}^{\text{new}}(\mathbf{x}) - \mathbf{w}^{\text{can}}(T^{\text{new}}(\mathbf{x}))\|_1, \quad (9)$$

where  $\mathcal{X}^{\text{new}}$  is the set of 3D points sampled within the human box under the novel human pose. Note that we fix the parameters of  $\mathbf{w}^{\text{can}}$  during training. In practice, we train neural skinning fields under multiple novel human poses simultaneously. This is implemented by conditioning  $F_{\Delta\mathbf{w}}$  on multiple latent codes. With the deformation field  $T^{\text{new}}$ , our method uses equations (1) and (2) to produce the neural radiance field under the novel human pose.

**3D shape generation.** In addition to synthesizing images under novel human poses, our approach can also explicitly animate a reconstructed human mesh, similar to the traditional animation methods. In particular, we first discretize the human bounding box at the canonical space with a voxel size of  $5\text{mm} \times 5\text{mm} \times 5\text{mm}$  and evaluate the volume densities for all voxels, which are used to extract the human mesh with the Marching Cubes algorithm [37]. Then, blend weights of mesh vertices are inferred from the neural blend weight field  $\mathbf{w}^{\text{can}}$ . Finally, given a novel human pose, we use equation (3) to transform each vertex, resulting in a deformed mesh under the target pose. The reconstruction results are presented in the supplementary material.

## 4. Implementation details

The networks of our radiance field  $F_\sigma$  and  $F_c$  closely follow the original NeRF [41]. We only use the single-level NeRF and sample 64 points along each camera ray. The network of  $F_{\Delta\mathbf{w}}$  is almost the same as that of  $F_\sigma$ , except that the final output layer of  $F_{\Delta\mathbf{w}}$  has 24 channels. In addition,  $F_{\Delta\mathbf{w}}$  applies  $\exp(\cdot)$  to the output. The details of network architectures are described in the supplementary material. The appearance code  $\ell_i$  and blend weight field code  $\psi_i$  both have dimensions of 128.

**Training.** Our method takes a two-stage training pipeline. First, we train the parameters of  $F_\sigma$ ,  $F_c$ ,  $F_{\Delta\mathbf{w}}$ ,  $\{\ell_i\}$  and  $\{\psi_i\}$  jointly over the input video. Second, neural blend weight fields under novel human poses are learned using equation (9). The Adam optimizer [26] is adopted for the training. The learning rate starts from  $5e^{-4}$  and decays exponentially to  $5e^{-5}$  along the optimization. The training is conducted on four 2080 Ti GPUs. For a three-view video of 300 frames, the first stage training takes around 200k iterations to converge (about 12 hours). For 200 novel human poses, the second stage training takes around 10k iterations to converge (about 30 minutes).

## 5. Experiments

### 5.1. Dataset and metrics

**H36M [19]** records multi-view videos with 4 cameras and collects human poses using the marker-based motion capture system. It includes multiple subjects performing complex actions. We select representative actions, split the videos into training and test frames, and perform experiments on subjects S1, S5, S6, S7, S8, S9, and S11. Three cameras are used for training and the remaining camera is selected for test. We use [22] to obtain the SMPL parameters from the 3D human poses and apply [15] to segment foreground humans. More details of training and test data can be found in the supplementary material.

**ZJU-MoCap [49]** records multi-view videos with 21 cameras and collects human poses using the marker-less motion capture system. For evaluation, we select four representative sequences: “Twirl”, “Taichi”, “Warmup”, and “Punch1”. Four uniformly distributed cameras are used for training and the remaining cameras for testing. We follow the experimental protocol in [49].

**Metrics.** Following typical protocols [41], we evaluate our method on image synthesis using two metrics: peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM). For 3D reconstruction, since there is no ground-truth geometry, we only provide qualitative results, which can be found in the supplementary material.

### 5.2. Performance on image synthesis

**Baselines.** We compare with state-of-the-art image synthesis methods [60, 62, 49] that also utilize SMPL priors. 1) Neural Textures [60] renders a coarse mesh with latent texture maps and uses a 2D CNN to interpret feature maps into target images. Since [60] is not open-sourced, we reimplement it and take the SMPL mesh as the input mesh. 2) NHR [62] extracts 3D features from input point clouds and renders them into 2D feature maps, which are then transformed into images using 2D CNNs. Since dense point clouds are difficult to obtain from sparse camera views, we take SMPL vertices as input point clouds. 3) Neural body [49] represents the human body with an implicit field conditioned on the latent codes anchored on the vertices of SMPL and renders the images using volume rendering.

**Results of novel view synthesis.** For comparison, we synthesize novel views of training video frames. Table 1 shows the comparison of our method with [60, 62]. Specifically, our model outperforms [60, 62] by a margin of at least 2.07 in terms of the PSNR metric and 0.024 in terms of the SSIM metric. Moreover, the proposed method achieves comparable results with the most recent state-of-the-art approach [49] as shown in Table 2, despite not being specifically designed for the novel view synthesis task.

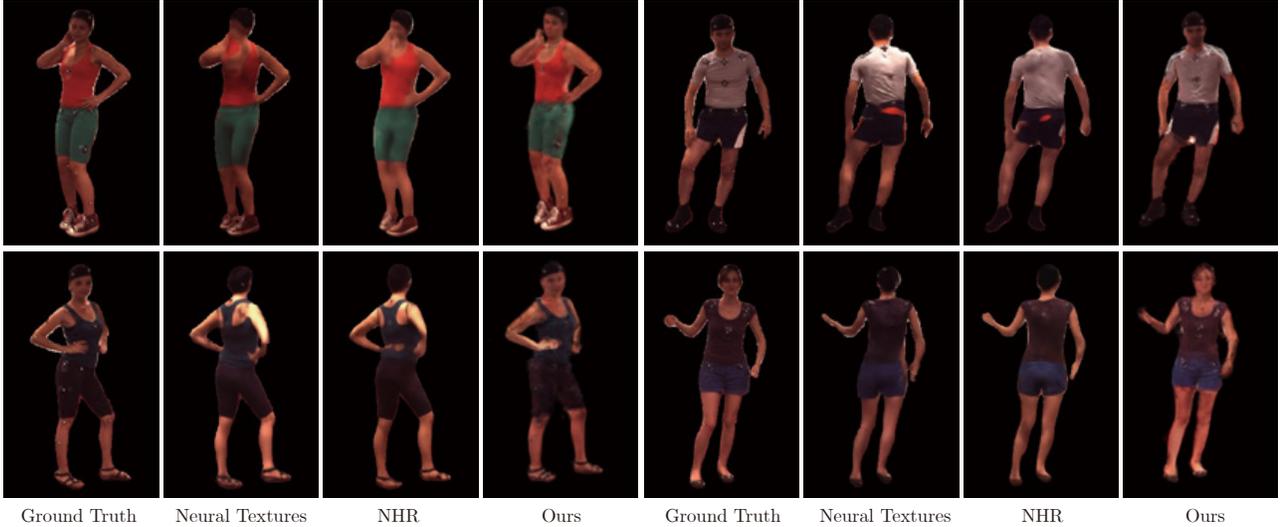


Figure 3: **Qualitative results of novel view synthesis on the H36M dataset.** [60, 62] have difficulty in controlling the viewpoint and seem to overfit training views. Compared with them, our method accurately renders the target view.

	PSNR			SSIM		
	NT [60]	NHR [62]	Ours	NT [60]	NHR [62]	Ours
S1	20.98	21.08	<b>22.05</b>	0.860	0.872	<b>0.888</b>
S5	19.87	20.64	<b>23.27</b>	0.855	0.872	<b>0.892</b>
S6	20.18	20.40	<b>21.13</b>	0.816	0.830	<b>0.854</b>
S7	20.47	20.29	<b>22.50</b>	0.856	0.868	<b>0.890</b>
S8	16.77	19.13	<b>22.75</b>	0.837	0.871	<b>0.898</b>
S9	22.96	23.04	<b>24.72</b>	0.873	0.879	<b>0.908</b>
S11	21.71	21.91	<b>24.55</b>	0.859	0.871	<b>0.902</b>
average	20.42	20.93	<b>23.00</b>	0.851	0.866	<b>0.890</b>

Table 1: **Results of novel view synthesis on H36M dataset in terms of PSNR and SSIM (higher is better).** “NT” means Neural Textures.

	PSNR				SSIM			
	NT [60]	NHR [62]	NB [49]	Ours	NT [60]	NHR [62]	NB [49]	Ours
novel view	22.61	23.25	<b>28.90</b>	27.10	0.899	0.905	<b>0.967</b>	0.949
novel pose	21.55	21.88	23.06	<b>23.16</b>	0.860	0.863	0.879	<b>0.893</b>

Table 2: **Results of novel view synthesis and novel pose synthesis on ZJU-MoCap dataset in terms of PSNR and SSIM (higher is better).** “NB” means Neural Body.

Figure 3 presents the qualitative comparison of our method with [60, 62]. Both [60, 62] have difficulty in controlling the rendering viewpoint and tend to synthesize contents of training views. As shown in the second person of Figure 3, they render the human back that is seen during training. In contrast, our method is able to accurately control the viewpoint, thanks to the explicit 3D representation.

**Results of novel pose synthesis.** For comparison, we synthesize test video frames from the test camera view. Table 3 compares our method with [60, 62] in terms of the PSNR metric and the SSIM metric. For both metrics, our method gives the best performances. Table 2 shows that our model also outperforms [49] when generating images under novel human poses on ZJU-MoCap dataset.

	PSNR			SSIM		
	NT [60]	NHR [62]	Ours	NT [60]	NHR [62]	Ours
S1	20.09	20.48	<b>21.37</b>	0.837	0.853	<b>0.868</b>
S5	20.03	20.72	<b>22.29</b>	0.843	0.860	<b>0.875</b>
S6	20.42	20.47	<b>22.59</b>	0.844	0.856	<b>0.884</b>
S7	20.03	19.66	<b>22.22</b>	0.838	0.852	<b>0.878</b>
S8	16.69	18.83	<b>21.78</b>	0.824	0.855	<b>0.882</b>
S9	22.20	22.18	<b>23.72</b>	0.851	0.860	<b>0.886</b>
S11	21.72	22.12	<b>23.91</b>	0.854	0.867	<b>0.889</b>
average	20.17	20.64	<b>22.55</b>	0.842	0.858	<b>0.880</b>

Table 3: **Results of novel pose synthesis on H36M dataset in terms of PSNR and SSIM (higher is better).** “NT” means Neural Textures.

The qualitative results are shown in Figure 4. For complex human poses, [60, 62, 49] give blurry and distorted rendering results. In contrast, synthesized images of our method achieve better visual quality. The results indicate that our model has better controllability on the image generation process than CNN-based methods.

### 5.3. Ablation studies

We conduct ablation studies on one subject (S9) of the H36M [19] dataset in terms of the novel pose synthesis performance. First, to analyze the benefit of learning  $F_{\Delta w}$ , we compare neural blend weight field with SMPL blend weight field. Then, to explore the influence of human pose accuracy, we estimate SMPL parameters from predicted human poses [7, 22] and perform training on these parameters. Finally, we explore the performances of our method under different numbers of video frames and camera views. Tables 4, 5, 6, and 7 summarize the results of ablation studies.

**Impact of neural blend weight field.** Table 4 shows the quantitative comparisons, which indicate that neural blend weight field performs better than SMPL blend weight field.

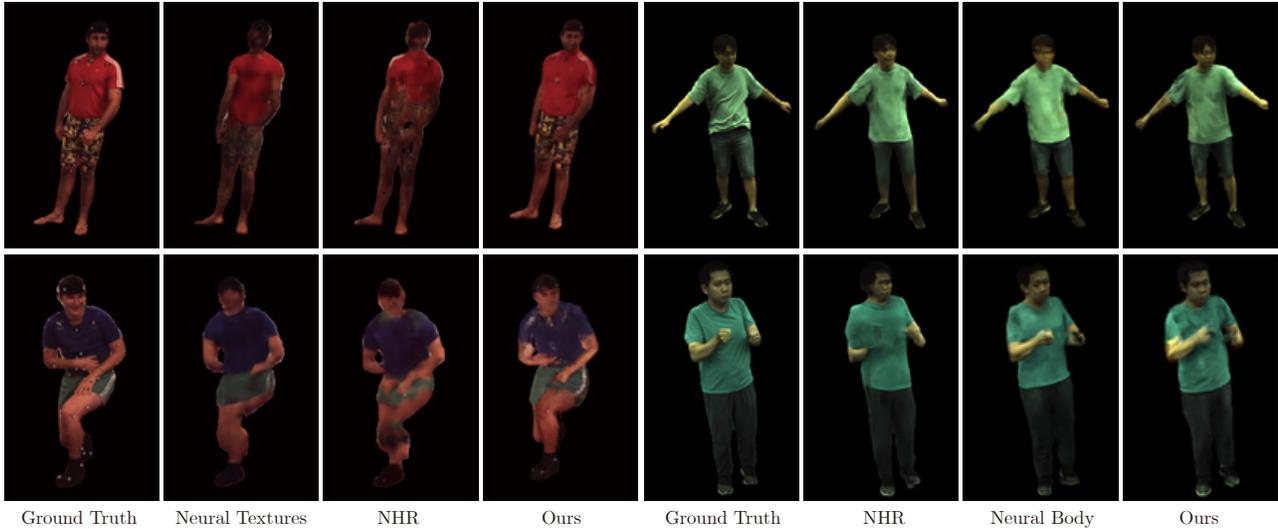


Figure 4: **Qualitative results of novel pose synthesis on the H36M and ZJU-MoCap datasets.** For complex human poses, [60, 62, 49] tend to generate distorted rendering results. In contrast to them, our method has a better generalization ability.

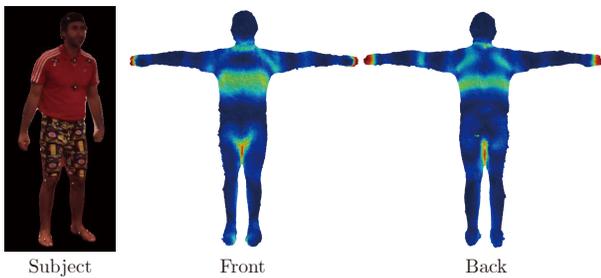


Figure 5: **Visualization of the residual vector field  $F_{\Delta w}$**  on the reconstructed geometries of subjects “S9” and “S6”. Red means large residual. Best viewed in color.

	PSNR	SSIM
Neural blend weight field	<b>23.72</b>	<b>0.886</b>
SMPL blend weight field	21.65	0.850

Table 4: **Comparison between neural blend weight field and SMPL blend weight field** on subject “S9”.

To better show the improvement on the SMPL blend weight field, Figure 5 visualizes the residual vector field  $F_{\Delta w}$  on our reconstructed geometry at the canonical space. The bigger residual has a redder color. We can see that regions of big residual mainly locate on the neck, hand, chest, and pants, which are human-specific details that SMPL cannot describe. The results indicate that our learned  $F_{\Delta w}$  are physically interpretable.

**Impact of the human pose accuracy.** Table 5 compares the models trained with human poses from marker-based and marker-less systems. The results show that more accurate human poses produce better rendering quality. The qualitative comparison is presented in Figure 6.

	PSNR	SSIM
Marker-based pose estimation	<b>23.72</b>	<b>0.886</b>
Marker-less pose estimation	22.27	0.858

Table 5: **Comparison between models trained with human poses** from marker-based and marker-less pose estimation methods on subject “S9”.

Frames	1	100	200	800
PSNR	20.29	23.40	<b>23.69</b>	23.16
SSIM	0.849	0.881	<b>0.883</b>	0.875

Table 6: **Results of models trained with different numbers of video frames** on subject “S9” of H36M dataset.

**Impact of the video length.** For comparison, we take 1, 100, 200 and 800 video frames for training and test the models on the same motion sequence. Table 6 lists the quantitative results of our models trained with different numbers of video frames. The results demonstrate that training on the video helps the representation learning, but the network seems to have difficulty in fitting very long videos. Empirically, we find that 150~300 frames are suitable for most subjects. Figure 7 presents the qualitative comparisons.

**Impact of the number of input views.** For comparison, we take one view for test and select 1, 2, and 3 nearest views for training. Table 7 compares the performances of models trained with different numbers of input views. Surprisingly, the three models have similar quantitative performances. Figure 8 further compares the three models, which shows that the model trained on 3 views renders more details. It is worth noting that the model trained on a single view already achieves reasonable rendering quality.

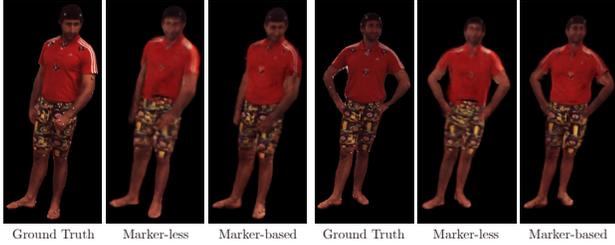


Figure 6: **Qualitative results of models trained on poses from marker-less and marker-based systems.**

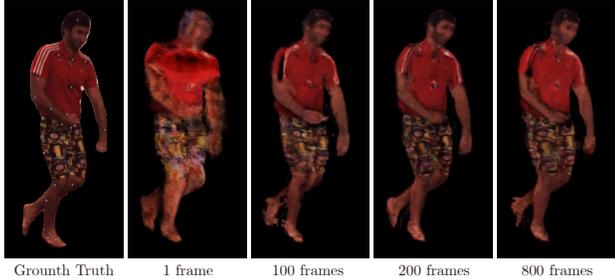


Figure 7: **Comparison of models trained with different numbers of video frames on the subject ‘S9’.**

	1 view	2 views	3 views
PSNR	23.81	<b>24.16</b>	23.72
SSIM	0.877	0.880	<b>0.886</b>

Table 7: **Results of models trained with different numbers of camera views on subject ‘S9’.**

#### 5.4. Running time

For  $512 \times 512$  images, our algorithm takes 1.09s to render an image on a desktop with an Intel i7 3.7GHz CPU and a GTX 1080 Ti GPU. Specifically, our implementation takes 0.39s for predicting the color and density fields, 0.63s for predicting the blend weight fields, and 0.07s for volume rendering. Because the number of points sampled along the ray is only 64 and the scene bound of a human is small, the rendering speed of our method is relatively fast.

#### 6. Limitations

Combining neural radiance fields with blend weight fields enables us to obtain impressive performances on novel view synthesis and novel pose synthesis. However, our method has a few limitations. 1) The skeleton-driven deformation model [29] cannot express the complex non-rigid deformations of garments. As a result, the performance of our method tends to degrade when reconstructing performers that wear loose clothes. It would be interesting to augment neural radiance fields with the deformation graph [43] that can model local garment deformations. 2) Currently our method requires rather accurate 3D human

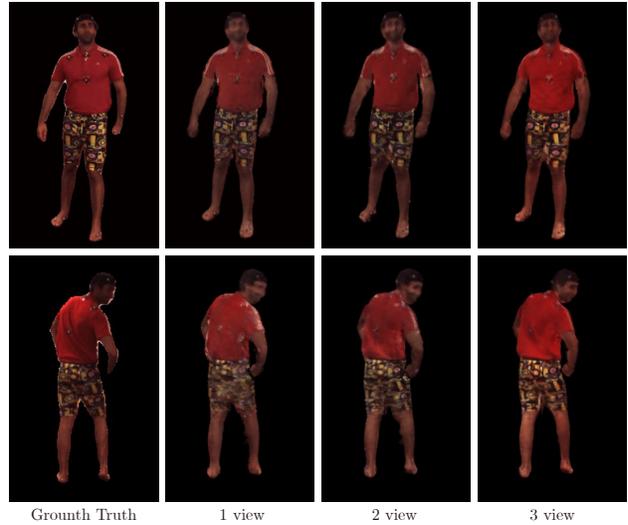


Figure 8: **Comparison of models trained with different numbers of camera views on the subject ‘S9’.**

skeletons. We hope that, in the future, we can find a way to refine human poses during training. 3) Same to NeRF, our proposed model is trained per-scene, which requires a lot of time to produce animatable human models. Generalizing the networks across different videos and reducing training time is left as future work. 4) Moreover, the rendering time of our model is a bit high. It could be solved with recent caching-based techniques [66, 17].

#### 7. Conclusion

We introduced a novel dynamic human representation for modeling animatable human characters from multi-view videos. Our method augments a neural radiance field with deformation fields that transform observation-space points to the canonical space. The deformation fields are constructed based on the skeleton-driven deformation framework, where we learn neural blend weight fields to generate observation-to-canonical and canonical-to-observation correspondences. The animatable neural radiance field is learned over the multi-view video with volume rendering and the consistency among blend weight fields. After training, our method can synthesize free-viewpoint videos of a performer given novel motion sequences. Experiments on the H36M and ZJU-MoCap datasets demonstrated that the proposed model achieves state-of-the-art performances on image synthesis under novel views and novel human poses.

**Acknowledgements:** The authors from Zhejiang University would like to acknowledge the support from the National Key Research and Development Program of China (No. 2020AAA0108901) and NSFC (No. 62172364).

## References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *ECCV*, 2020. 2
- [2] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single RGB camera. In *CVPR*, 2019. 2
- [3] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, 2018. 2
- [4] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *ICCV*, 2019. 2
- [5] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Combining implicit function learning and parametric models for 3d human reconstruction. In *ECCV*, 2020. 2
- [6] 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. In *NeurIPS*, 2020. 4
- [7] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: realtime multi-person 2d pose estimation using part affinity fields. *PAMI*, 2018. 6
- [8] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *ICCV*, 2019. 2
- [9] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *CVPR*, 2020. 2
- [10] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM TOG*, 2015. 1, 2
- [11] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *SIGGRAPH*, 2000. 2
- [12] Boyang Deng, JP Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Nasa: Neural articulated shape approximation. In *ECCV*, 2020. 2
- [13] Junting Dong, Qing Shuai, Yuanqing Zhang, Xian Liu, Xiaowei Zhou, and Hujun Bao. Motion capture from internet videos. In *ECCV*, 2020. 2
- [14] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4d: Real-time performance capture of challenging scenes. *ACM TOG*, 2016. 1, 2
- [15] Ke Gong, Xiaodan Liang, Yicheng Li, Yimin Chen, Ming Yang, and Liang Lin. Instance-level human parsing via part grouping network. In *ECCV*, 2018. 3, 5
- [16] 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 TOG*, 2019. 1, 2
- [17] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *ICCV*, 2021. 8
- [18] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *CVPR*, 2020. 2, 4
- [19] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *PAMI*, 2013. 2, 3, 5, 6
- [20] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 2
- [21] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *CVPR*, 2020. 2
- [22] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *CVPR*, 2018. 2, 3, 4, 5, 6
- [23] James T Kajiya. The rendering equation. In *SIGGRAPH*, 1986. 4
- [24] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 2
- [25] Ladislav Kavan, Steven Collins, Jiří Žára, and Carol O’Sullivan. Skinning with dual quaternions. In *I3D*, 2007. 2
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [27] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019. 2
- [28] 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 *ECCV*, 2020. 2
- [29] John P Lewis, Matt Cordner, and Nickson Fong. Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In *SIGGRAPH*, 2000. 1, 2, 3, 8
- [30] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 2, 3
- [31] Zhengqi Li, Wenqi Xian, Abe Davis, and Noah Snavely. Crowdsampling the plenoptic function. In *ECCV*, 2020. 2
- [32] Yariv Lior, Kasten Yoni, Moran Dror, Galun Meirav, Atzmon Matan, Basri Ronen, and Lipman Yaron. Multiview neural surface reconstruction by disentangling geometry and appearance. In *NeurIPS*, 2020. 2
- [33] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *NeurIPS*, 2020. 2
- [34] Lingjie Liu, Weipeng Xu, Marc Habermann, Michael Zollhoefer, Florian Bernard, Hyeonwoo Kim, Wenping Wang, and Christian Theobalt. Neural human video rendering by learning dynamic textures and rendering-to-video translation. *TVCG*, 2020. 2

- [35] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. In *SIGGRAPH*, 2019. 2
- [36] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM TOG*, 2015. 2, 4
- [37] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH*, 1987. 5
- [38] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *NeurIPS*, 2017. 2
- [39] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable person image synthesis with attribute-decomposed gan. In *CVPR*, 2020. 2
- [40] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019. 2
- [41] 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 *ECCV*, 2020. 1, 2, 3, 4, 5
- [42] Ryota Natsume, Shunsuke Saito, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, and Shigeo Morishima. Siclope: Silhouette-based clothed people. In *CVPR*, 2019. 2
- [43] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *CVPR*, 2015. 8
- [44] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 2020. 2
- [45] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 2
- [46] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin Brualla. Deformable neural radiance fields. *arXiv preprint arXiv:2011.12948*, 2020. 1, 2, 3
- [47] 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 *CVPR*, 2019. 4
- [48] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3d human pose and shape from a single color image. In *CVPR*, 2018. 2
- [49] 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 *CVPR*, 2021. 2, 5, 6, 7
- [50] Sergey Prokudin, Michael J Black, and Javier Romero. Smpix: Neural avatars from 3d human models. In *WCCV*, 2021. 2
- [51] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 1, 2, 3
- [52] Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. Anr: Articulated neural rendering for virtual avatars. In *CVPR*, 2021. 2
- [53] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM ToG*, 2017. 4
- [54] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 2
- [55] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, 2020. 2
- [56] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 1
- [57] Aliaksandra Shysheya, Egor Zakharov, Kara-Ali Aliiev, Renat Bashirov, Egor Burkov, Karim Isakov, Aleksei Ivakhnenko, Yury Malkov, Igor Pasechnik, Dmitry Ulyanov, et al. Textured neural avatars. In *CVPR*, 2019. 2
- [58] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *NeurIPS*, 2019. 2
- [59] Zhuo Su, Lan Xu, Zerong Zheng, Tao Yu, Yebin Liu, et al. Robustfusion: Human volumetric capture with data-driven visual cues using a rgb camera. In *ECCV*, 2020. 2
- [60] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM TOG*, 2019. 2, 5, 6, 7
- [61] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Photo wake-up: 3d character animation from a single photo. In *CVPR*, 2019. 2
- [62] Minye Wu, Yuehao Wang, Qiang Hu, and Jingyi Yu. Multi-view neural human rendering. In *CVPR*, 2020. 2, 5, 6, 7
- [63] 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 *CVPR*, 2020. 4
- [64] 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 *CVPR*, 2021. 2
- [65] Jae Shin Yoon, Lingjie Liu, Vladislav Golyanik, Kripasindhu Sarkar, Hyun Soo Park, and Christian Theobalt. Pose-guided human animation from a single image in the wild. In *CVPR*, 2021. 2
- [66] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. 8
- [67] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. In *ICCV*, 2019. 2