

# SurMo: Surface-based 4D Motion Modeling for Dynamic Human Rendering

Tao Hu, Fangzhou Hong, Ziwei Liu  
 S-Lab, Nanyang Technological University, Singapore

## Abstract

Dynamic human rendering from video sequences has achieved remarkable progress by formulating the rendering as a mapping from static poses to human images. However, existing methods focus on the human appearance reconstruction of every single frame while the temporal motion relations are not fully explored. In this paper, we propose a new 4D motion modeling paradigm, **SurMo**, that jointly models the temporal dynamics and human appearances in a unified framework with three key designs: **1) Surface-based motion encoding** that models 4D human motions with an efficient compact surface-based triplane. It encodes both spatial and temporal motion relations on the dense surface manifold of a statistical body template, which inherits body topology priors for generalizable novel view synthesis with sparse training observations. **2) Physical motion decoding** that is designed to encourage physical motion learning by decoding the motion triplane features at timestep  $t$  to predict both spatial derivatives and temporal derivatives at the next timestep  $t + 1$  in the training stage. **3) 4D appearance decoding** that renders the motion triplanes into images by an efficient volumetric surface-conditioned renderer that focuses on the rendering of body surfaces with motion learning conditioning. Extensive experiments validate the state-of-the-art performance of our new paradigm and illustrate the expressiveness of surface-based motion triplanes for rendering high-fidelity view-consistent humans with fast motions and even motion-dependent shadows. Our project page is at: <https://taohuimd.github.io/projects/SurMo>.

## 1. Introduction

Creating volumetric videos of human actors is required in many applications such as AR/VR, telepresence, video game and film character creation. Recent neural rendering methods [25, 43, 54, 61] have made great progress in generating free-viewpoint videos of humans from several sparse multi-view videos, which are simple yet effective compared with traditional graphics approaches [2, 3, 62]. At the core of the technique is to model pose- and time-varying appearances of dynamic humans. The motions of dressed humans are often expressed as the movement of body and a sequence of natural secondary motion of clothes, e.g., dy-




Figure 1. Given several sparse multi-view video sequences with estimated 3D body meshes, SurMo synthesizes subject-specific appearance. We specifically focus on the synthesis of plausible time-varying appearances by learning an effective 4D motion representation.

dynamic movements of the T-shirt induced by dancing in Fig. 1. The secondary motion of clothes arises from intricate physical interactions with the body, typically changing over time, which leads to challenges for plausible rendering of dynamic humans.

Most existing methods [25, 43, 61] are conditioned on static poses and use a pose-guided generator to synthesize time-varying appearances. However, the appearance of clothed humans undergoes complex geometric transformations induced not only by the static pose but also its dynamics, whereas the modeling of dynamics is often ignored. In addition, these methods are often focused on the 2D appearance reconstruction of every single frame while the temporal motion relations are not fully explored. Due to these issues, they fail to render plausible secondary motions of clothes, *i.e.*, generating the same appearances for fast and slow motions. A key challenge of learning temporal dynamics lies in the requirement of a tremendous amount of training observations to construct a 4D motion volume.

To solve these issues, we propose a new paradigm to learn time-varying appearances of the secondary motion from just several sparse viewpoint video sequences, which is achieved by jointly modeling the temporal motion dynamics and human appearances in a unified rendering framework based on an efficient surface-based motion

representation. At the core of the paradigm is a feature encoder-decoder framework with three key components: **1)** surface-based motion encoding; **2)** physical motion decoding; and **3)** 4D appearance decoding.

Firstly, in contrast to existing pose-guided methods [25, 42, 43, 61] that focus on static poses as a conditional variable, we extract an expressive 4D motion input from the 3D body mesh sequences obtained from training video as our input, which includes both a static pose represented by a spatial 3D mesh and its temporal dynamics. Furthermore, we notice that the non-rigid deformations of garments typically occur around the body surface instead of in a 3D volume, and hence we propose to model human motions on the body surface by projecting the extracted spatial-temporal 4D motion input to the dense 2D surface UV manifold of a clothless body template (*e.g.*, SMPL [28]). To model temporal clothing offsets, a motion encoder is employed to lift the clothless motion features into a motion triplane that encodes both spatial and temporal motion relations in a compact 3D triplane with time-varying dynamics conditioned. The triplane is defined in the surface  $\mathbf{u} - \mathbf{v} - \mathbf{h}$  coordinate system, with  $\mathbf{u} - \mathbf{v}$  to represent the motion of the clothless body template, and an extensional coordinate  $\mathbf{h}$  to represent the secondary motion of clothes, *i.e.*, the temporal clothing offsets are parameterized by a signed distance to the body surface. In this way, 4D motions can be effectively represented by a surface-based triplane.

Secondly, we propose to physically model the spatial and temporal motion dynamics in the rendering network. Specifically, with the surface-based triplane conditioned at time  $t$ , a motion decoder is employed to decode the motion triplanes to predict the motion at the next timestep  $t+1$ , *i.e.*, spatial derivative of the motion physically corresponding to surface normal map and temporal derivatives corresponding to surface velocity. We illustrate the physical motion learning significantly improves the rendering quality.

Thirdly, the motion triplanes are decoded into high-quality images at two stages: a volumetric surface-conditioned renderer that is focused on the rendering around human body surface and filters the query points far from the body surface for efficient volumetric rendering, and a geometry-aware super-resolution module for efficient high-quality image synthesis.

We conduct a systematical analysis of how human appearances are affected by temporal dynamics, and it was observed that some baseline methods mainly generate pose-dependent appearances instead of time-varying appearances for free-view video generations. In addition, quantitative and qualitative experiments are performed on three datasets with a total of 9 subject sequences, including ZJU-MoCap [43], AIST++ [24], and MPII-RDDC [13], which validate the effectiveness of SurMo in different scenarios.

In summary, our contributions are:

**1)** A new paradigm for learning dynamic humans from videos that jointly models temporal motions and human appearances in a unified framework, and one of the early works that systematically analyzes how human appearances are affected by temporal dynamics.

**2)** An efficient surface-based triplane that encodes both spatial and temporal motion relations for expressive 4D motion modeling.

**3)** We achieve state-of-the-art results and show that our new paradigm is capable of learning high-fidelity appearances from fast motion sequences (*e.g.*, AIST++ dance videos) or synthesizing motion-dependent shadows in challenging scenarios.

## 2. Related Work

Our method is closely related to many sub-fields of visual computing, and below we discuss a set of the work.

**3D Shape Representations.** To capture detailed shapes of 3D objects, most recent papers utilize implicit functions [7, 8, 20, 33–35, 40, 41, 48–50, 56, 57, 65, 66] or point clouds [16, 31, 32] due to their topological flexibility. These methods aim to learn geometry from 3D datasets, whereas we synthesize human images of novel poses only from 2D RGB training images.

**Rendering Humans by 2D GANs.** Some existing approaches render human avatars by neural image translation, *i.e.*, they utilize GAN [10] networks to learn a mapping from poses (given in the form of renderings of a skeleton [4, 23, 45, 52, 59, 68], dense mesh [11, 26, 27, 36, 51, 59] or joint position heatmaps [1, 29, 30]) to human images [4, 59]. To improve temporal stability, some methods [17, 44, 46, 55] propose to utilize the SMPL [28] priors for pose-guided generations. However, these methods do not reconstruct geometry explicitly and cannot handle self-occlusions effectively.

**Rendering Humans by 3D-aware Renderer.** For stable view synthesis, recent papers [6, 25, 38, 42, 43, 53] propose to unify geometry reconstruction with view synthesis by volume rendering, which, however, is computationally heavy. To solve this issue, some recent 3D-GAN papers [5, 12, 15, 18, 19, 37, 39, 67] propose a hybrid rendering strategy for efficient geometry-aware rendering, that is, render low-resolution volumetric features for geometry learning, and employ a super-resolution module for high-resolution image synthesis. We adopt this strategy for efficient rendering, whereas ours is distinguished by rendering articulated humans with 4D motion modeling.

**UV-based Pose Representation.** Some methods [18, 19, 31, 32, 47, 63] propose to project 3D posed meshes into a 2D positional map for pose encoding, which can be used for different downstream tasks, such as 3D reconstruction and novel view synthesis. [31, 32] rely on 3D supervision to learn 3D reconstructions, and they do not take as input the

dynamics. To improve the rendering quality, [18, 47] proposes to utilize additional driving views as input for faithful rendering in telepresence applications. However, they do not explore how to learn temporal dynamics from pose sequences. By taking as input a 3D normal and velocity map, [63] encodes motion dynamics in their rendering network, whereas they do explicitly learn motions, *i.e.*, predicting the motion status at the next timestep, and besides, [63] is a 2D rendering method, and they do not explicitly learn motions.

### 3. Methodology

#### 3.1. Problem Setup

Given a sparse multi-view video of a clothed human in motion and corresponding 3D pose estimations  $\{\mathbf{P}_0, \dots, \mathbf{P}_t\}$ , our goal is to synthesize time-varying appearances of the individual under novel views. Existing methods [25, 42, 43, 61] formulate the problem of human rendering as learning a representation via a feature encoder-decoder framework:

$$\begin{aligned} \mathbf{f}_t &= \mathcal{E}_{\mathcal{P}}(\mathbf{P}_t, \mathbf{z}_t) \\ \mathbf{I}_t &= \mathcal{D}_{\mathcal{R}}(\mathbf{f}_t) \end{aligned} \quad (1)$$

where a pose encoder  $\mathcal{E}_{\mathcal{P}}$  takes as input a representation of posed body  $\mathbf{P}_t$  (*e.g.*, 2D or 3D keypoints or 3D body mesh vertices) and a timestamp embedding  $\mathbf{z}_t$  at time  $t$ , and outputs intermediate pose-dependent features  $\mathbf{f}_t$  that can be rendered by a decoder  $\mathcal{D}_{\mathcal{R}}$  to reconstruct the appearance images  $\mathbf{I}_t$  of the corresponding pose at time  $t$ . The time embedding  $\mathbf{z}_t$  often serves as a residual that is updated passively in backpropagation by image reconstruction loss.

However, there are two issues with the above well-adopted paradigm. First, the appearance of clothed humans undergoes complex geometric transformations induced not only by the static pose  $\mathbf{P}_t$  but also its dynamics, whereas the modeling of dynamics is ignored in E.q. 1, and the residual  $\mathbf{z}_t$  cannot expressively model physical dynamics neither. Second, existing methods focus on the per-image reconstruction while ignoring the temporal relations of a motion sequence in training, *i.e.*, sampling input pose  $\mathbf{P}_t$  and timestamp  $\mathbf{z}_t$  individually, and supervising image reconstructions frame-by-frame, partly because defining temporal supervision in the 2D image space is challenging, especially for articulated humans that often suffer from misalignment problems due to pose estimation errors.

To solve these issues, we propose a new paradigm for learning view synthesis of dynamic humans from video sequences, which is formulated as:

$$\begin{aligned} \mathbf{f}_t &= \mathcal{E}_{\mathcal{M}}(\mathbf{P}_t, \mathbf{D}_t) \\ \frac{\partial \mathbf{P}_{t+1}}{\partial \mathbf{x}}, \frac{\partial \mathbf{P}_{t+1}}{\partial t} &= \mathcal{D}_{\mathcal{M}}(\mathbf{f}_t) \\ \mathbf{I}_t &= \mathcal{D}_{\mathcal{R}}(\mathbf{f}_t) \end{aligned} \quad (2)$$

where the input is represented by dynamic motions including a static pose  $\mathbf{P}_t$  and its physical dynamics  $\mathbf{D}_t$  at time  $t$ , and a motion encoder  $\mathcal{E}_{\mathcal{M}}$  is employed to encode the inputs into motion-dependent features (Sec. 3.2). Besides, in the training stage, a physical motion decoder  $\mathcal{D}_{\mathcal{M}}$  is introduced to enforce the learning of spatial and temporal relations by decoding the intermediate feature  $\mathbf{f}_t$  to predict the spatial derivatives at  $t+1$  physically corresponding to surface normal  $\mathbf{N}_{t+1} = \frac{\partial \mathbf{P}_{t+1}}{\partial \mathbf{x}}$ , and temporal derivatives at  $t+1$  corresponding to surface velocity  $\mathbf{V}_{t+1} = \frac{\partial \mathbf{P}_{t+1}}{\partial t}$  (Sec. 3.3).  $\mathbf{f}_t$  is rendered into human images by a decoder  $\mathcal{D}_{\mathcal{R}}: \mathcal{G}_2 \circ \mathcal{G}_1$  (Sec. 3.5). The framework is depicted in Fig. 2.

**Notation.** In the following parts, we use  $\mathbf{f}_t^{\text{space}}$  to denote the features in the pipeline, *i.e.*,  $\mathbf{f}_t^{\text{3D}}$  is a 4D representation that is defined in 3D space with temporal  $t$  conditioned, and  $\mathbf{M}_t^{\text{3D}}$  is 4D motion input defined in 3D space with  $t$  conditioned.

#### 3.2. Surface-based 4D Motion Encoding

**Extracting 4D Motions.** We first extract 4D motions from a sequence of time-varying parametric posed body (*e.g.*, SMPL) meshes obtained from training video sequences. We describe the motion  $\mathbf{M}_t^{\text{3D}}$  at time  $t$  as a static skeleton pose  $\mathbf{P}_t$  and its dynamics  $\mathbf{D}_t$ . The dynamics at  $t$  are physically determined by the current pose, 3D velocity, and motion trajectory of the past several timesteps, which also contribute to the time-varying appearances of the secondary motion. The pose  $\mathbf{P}_t$  is represented by the 3D vertices of the posed mesh, and dynamics  $\mathbf{D}_t$  are parameterized by 1) body surface velocity  $\mathbf{V}_t$  corresponding to the temporal derivatives of the current pose at  $t$ , and 2) motion trajectory  $\mathbf{T}_t$  that aggregates the temporal derivatives over the past several timesteps with a sliding window size of  $w$  with weight  $\lambda$ :

$$\begin{aligned} \mathbf{M}_t^{\text{3D}} &= [\mathbf{P}_t, \mathbf{D}_t], \quad \mathbf{D}_t = [\mathbf{V}_t, \mathbf{T}_t] \\ \mathbf{V}_t &= \frac{\partial \mathbf{P}_t}{\partial t}, \quad \mathbf{T}_t = \mathbf{P}_t + \frac{1}{\sum \lambda_i} \sum_{i=1}^w \mathbf{V}_{t-i} * \lambda_i \end{aligned} \quad (3)$$

Note the motion trajectory aggregated from several consecutive timesteps makes the motion representation robust to pose estimation errors, *i.e.*, the pose estimations for two consecutive timesteps may be the same due to pose estimation errors in practice. An ablation study of the motion trajectory representation can be found in the supp. mat.

**Recording 4D Motions on the Surface Manifold.** Modeling the motions with temporal dynamics often requires dense observations to construct a 4D motion volume. Instead, we notice that non-rigid deformations of human geometry in motion typically occur around the body surface instead of a 3D volume, and hence we propose to model the motions on the human body surface. Specifically, we project the 4D motion input  $\mathbf{M}_t^{\text{3D}}$  including spatial pose

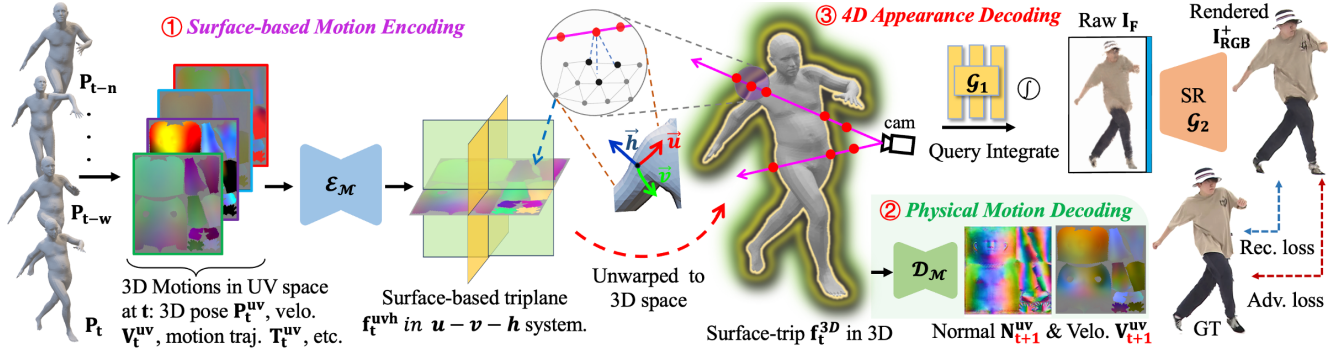


Figure 2. Framework overview. Given a set of time-varying 3D body meshes  $\{\mathbf{P}_t, \dots, \mathbf{P}_{t-n}\}$  obtained from training video sequences, we aim to synthesize high-fidelity appearances of a clothed human in motion via a feature encoder-decoder framework: *Motion Encoding*, and joint *Motion* and *Appearance Decoding*. **1)** We take as input an expressive 4D motion representation extracted from the mesh sequences including 3D pose, 3D velocity at time  $t$ , and motion trajectory over the past  $w$  timesteps that encode both spatial and temporal relations of the motion sequence, which are projected to the spatially aligned UV surface space. A motion encoder  $\mathcal{E}_M$  is employed to lift the 2D UV-aligned features to a 3D surface-based triplane  $\mathbf{f}_t^{\text{uvh}}$  in an UV-plus-height space with a signed distance height to model temporal clothing offsets. **2)** A motion decoder  $\mathcal{D}_M$  is designed to encourage physical motion learning in training by decoding the triplane features  $\mathbf{f}_t^{\text{uvh}}$  to predict the motion at the next timestep  $t+1$ , *i.e.* spatial derivatives surface normal  $\mathbf{N}_{t+1}^{\text{uv}}$  and temporal derivatives surface velocity  $\mathbf{V}_{t+1}^{\text{uv}}$  in UV space. **3)** Finally, given a target camera view, the triplane  $\mathbf{f}_t^{\text{uvh}}$  is rendered into high-quality images by a volumetric surface-conditioned renderer including volumetric low-resolution rendering by  $\mathcal{G}_1$  and an efficient geometry-aware super-resolution by  $\mathcal{G}_2$ .

and temporal dynamics from 3D space into a compact spatially aligned UV space using the geometric transformation  $\mathcal{W}$  that is pre-defined by the parametric (*e.g.*, SMPL) body template, which yields  $\mathbf{f}_t^{\text{uv}} = \mathcal{W}\mathbf{M}_t^{\text{3D}}$ . The UV-aligned motion feature  $\mathbf{f}_t^{\text{uv}}$  faithfully preserves the articulated structures of body topology in a compact 2D space.

**Generating Surface-based Triplanes for Motion Modeling.** To model clothed humans, we further employ a motion encoder  $\mathcal{E}_M$  that lifts the clothless 2D features  $\mathbf{f}_t^{\text{uv}}$  to a 3D triplane representation  $\mathbf{f}_t^{\text{uvh}}$  that is defined in a  $\mathbf{u} - \mathbf{v} - \mathbf{h}$  system, with  $\mathbf{u} - \mathbf{v}$  to represent the motion of the clothless body template, and an extensional coordinate  $\mathbf{h}$  to represent the secondary motion of clothes, *i.e.*, the temporal cloth offsets, and hence 4D clothed motions can be parameterized by a surface-based triplane by:

$$\mathbf{f}_t^{\text{uvh}} = \mathcal{E}_M(\mathbf{f}_t^{\text{uv}}) = \mathcal{E}_M(\mathcal{W}\mathbf{M}_t^{\text{3D}}), \quad (4)$$

where  $\mathcal{W}$  is the geometric transformation from 3D space to UV space.  $\mathbf{f}_t^{\text{uvh}}$  consists of three planes  $\mathbf{x}_{uv}, \mathbf{x}_{uh}, \mathbf{x}_{hv} \in \mathbb{R}^{U \times V \times C}$  to form the spatial relationship, where  $U, V$  denote spatial resolution and  $C$  is the channel number. In contrast to the volumetric triplanes used in EG3D [5] where the three planes represent three vertical planes in the 3D volumetric space,  $\mathbf{f}_t^{\text{uvh}}$  is defined on the human body surface inheriting human topology priors for motion modeling, as illustrated by an unwrapped triplane in 3D space in Fig 2.

### 3.3. Physical Motion Decoding

We propose a physical motion decoder  $\mathcal{D}_M$  to learn spatial and temporal features of motion, which is achieved by decoding the intermediate motion feature  $\mathbf{f}_t^{\text{uvh}}$  learned by  $\mathcal{E}_M$

to predict the motions at the next timestep  $t+1$ , such as spatial derivatives corresponding to surface normal  $\mathbf{N}_{t+1}^{\text{uv}}$ , temporal derivatives corresponding to surface velocity  $\mathbf{V}_{t+1}^{\text{uv}}$ . Note that we model the normal and velocity in the UV space by:

$$\{\mathbf{N}_{t+1}^{\text{uv}}, \mathbf{V}_{t+1}^{\text{uv}}\} = \mathcal{D}_M(\mathbf{f}_t^{\text{uvh}}) = \mathcal{D}_M \circ \mathcal{E}_M(\mathcal{W}\mathbf{M}_t^{\text{3D}}) \quad (5)$$

### 3.4. 4D Appearance Decoding

**Volumetric Surface-conditioned Rendering.** Given a target camera viewpoint, the triplane  $\mathbf{f}_t^{\text{uvh}}$  is first rendered into low-resolution volumetric features  $\mathbf{I}_F \in \mathbb{R}^{H \times W \times \hat{C}}$  by a conditional NeRF  $\mathbb{F}_\Phi$ , where  $H$  and  $W$  are the spatial resolution and  $\hat{C}$  is the channel number. To be more specific, given a 3D query point  $\mathbf{p}_i$ , we transform it into a surface-based local coordinate  $\hat{\mathbf{p}}_i = (\mathbf{u}_i, \mathbf{v}_i, \mathbf{h}_i)$  w.r.t. the tracked body mesh  $\mathbf{P}_t$ . Here we search the nearest face  $\mathbf{f}_i$  of  $\mathbf{P}_t$  for query point  $\mathbf{p}_i$  and  $(\mathbf{u}_i, \mathbf{v}_i)$  and  $\mathbf{h}_i$  are the barycentric coordinates of the nearest point on  $\mathbf{f}_i$ , and the signed distance respectively. We therefore obtain the local feature  $\mathbf{z}_i^{\text{uvh}}$  of the query point  $\mathbf{p}_i$  as

$$\mathbf{z}_i^{\text{uvh}} = \text{cat}[\Pi(\mathbf{x}_{uv}; \mathbf{u}_i, \mathbf{v}_i), \Pi(\mathbf{x}_{uh}; \mathbf{u}_i, \mathbf{h}_i), \Pi(\mathbf{x}_{hv}; \mathbf{h}_i, \mathbf{v}_i)], \quad (6)$$

where  $\Pi(\cdot)$  denotes sampling operation,  $\text{cat}[\cdot]$  denotes the concatenation operator.

Given camera direction  $\mathbf{d}_i$ , the appearance features  $\mathbf{c}_i$  and density features  $\sigma_i$  of point  $\mathbf{p}_i$  are predicted by

$$\{\mathbf{c}_i, \sigma_i\} = \mathbb{F}_\Phi(\mathbf{d}_i, \mathbf{z}_i^{\text{uvh}}) \quad (7)$$

We integrate all the radiance features of sampled points into a 2D feature map  $\mathbf{I}_F$  at time  $t$  through volume renderer




Figure 3. Qualitative comparisons on novel view synthesis on the subject S313 of ZJU-MoCap dataset. Two motion sequences S1 (swing arms left to right) and S2 (raise and lower arms) are shown. We specifically focus on the synthesis of time-varying appearances (especially T-shirt wrinkles), by evaluating the rendering results under similar poses yet with different movement directions, which are marked in the same color, such as the pairs of ①②, ③④, and ⑤⑥. Our method synthesizes high-fidelity time-varying appearances, whereas SOTA HumanNeRF generates almost the same cloth wrinkles.

$\mathcal{G}_1$  [21]

$$\mathbf{I}_{\mathbf{F}} = \mathcal{G}_1(\mathbf{f}_{\mathbf{t}}^{\text{uvh}}, \text{cam}; \mathbb{F}_{\Phi}) \quad (8)$$

where  $\mathbf{f}_{\mathbf{t}}^{\text{uvh}}$  is extracted from Eq. 4.

**Efficient Geometry-Aware Super-Resolution.** Sampling dense points to render the full-resolution volumetric features is computationally heavy. Instead, we employ a super-resolution network  $\mathcal{G}_2$  [5] to render high-resolution images  $\mathbf{I}_{\text{RGB}}^+ \in \mathbb{R}^{\hat{H} \times \hat{W} \times 3}$ :

$$\mathbf{I}_{\text{RGB}}^+ = \mathcal{G}_2 \circ \mathcal{G}_1(\mathbf{f}_{\mathbf{t}}^{\text{uvh}}, \text{cam}; \mathbb{F}_{\Phi}) \quad (9)$$

### 3.5. Optimization

SurMo is trained end-to-end to optimize  $\mathcal{E}_{\mathcal{M}}$ ,  $\mathcal{D}_{\mathcal{M}}$ , and renderers  $\mathcal{G}_1$ ,  $\mathcal{G}_2$  with 2D image loss. We employ Adversarial Loss and Reconstruction Loss including Pixel Loss,

Perceptual Loss, Face Identity Loss, Velocity and Normal Loss, and Volume Rendering Loss for supervision in training, with Adam [22] as the optimizer. Refer to the supp. mat. for more details.

## 4. Experiments

**Dataset and Metrics.** We evaluate the novel view synthesis on ZJU-MoCap, MPII-RDDC, and AIST++, with a total of 9 subjects. We use the same camera setup as Neural Body, that is, 4 cameras used for training, and the other 18 for testing. For MPII-RDDC, 18 cameras in training, and 9 for testing. Since ASIT++ only has 9 cameras, we use 6 for training, and the remaining 3 for testing. Refer to more details in the supp. mat. We compare each method on per-pixel metrics including SSIM [60] and PSNR, and




Figure 4. Qualitative comparisons on novel view synthesis on the subject S387, S315 of ZJU-MoCap dataset. Row 1 and 2 show similar poses occurring at different timesteps (not consecutive frames). The results indicate that our method synthesizes time-varying appearances while other methods mainly generate pose-dependent appearances.

Table 1. Quantitative comparisons on ZJU-MoCap dataset (averaged on all test views and poses on 6 sequences) for novel-view synthesis. To reduce the influence of the background, all scores are calculated from images cropped to 2D bounding boxes. Note that the perception metrics LPIPS [64] and FID [14] capture human judgment better than per-pixel metrics such as SSIM [60] or PSNR, as stated in [25, 51].

ZJU-S1-6	LPIPS ↓	FID ↓	SSIM ↑	PSNR ↑
Neural Body [43]	.164	125.10	.703	21.438
Instant-NVR [9]	.202	144.99	.738	21.748
HumanNeRF [61]	.106	88.382	.792	23.624
Ours	<b>.075</b>	<b>67.725</b>	<b>.833</b>	<b>24.815</b>

perception metrics including LPIPS [64] and FID [14].

**Baselines.** We compare our method against SOTA methods including Neural Body [43], HumanNeRF [61], Instant-NVR [9], ARAH [58], DVA [47] and HVTR++ [18]. Neural Body models human poses in 3D space with point cloud as pose representation, HumanNeRF and Instant-NVR model poses in a canonical space with inverse skinning, ARAH is a forward-skinning based method, DVA and HVTR++ encode poses in UV space, and take driving view signals as input for telepresence applications.

#### 4.1. Comparisons to SOTA Methods

**Quantitative comparisons** We conduct the quantitative comparisons on the three datasets with a total of 9 subjects. The quantitative results on ZJU-MoCap are summarized in Tab. 1, which suggests that our new paradigm outperforms the SOTA by a big margin, and we achieve the best quantitative results on all four metrics. The detailed compar-

isons on each subject of ZJU-MoCap, and quantitative results on MPII-RDDC and AIST++ are listed in the supp. mat. Refer to the comparisons with ARAH [58], DVA [47] and HVTR++ [18] in the supp. mat.

**Time-varying Appearances with Dynamics Conditioning on ZJU-MoCap.** We compare with baseline methods for novel view synthesis on two motion sequences of subject S313 from the ZJU-MoCap dataset, as shown in Fig. 3. We evaluate the capability of each method to synthesize time-varying appearances. Specifically, two sequences S1 (swing arms left to right) and S2 (raise and lower arms) are evaluated, where similar poses occur at different timesteps with different dynamics (*e.g.*, movement directions, trajectory, or velocity) are marked in the same color, such as the pairs of ①②, ③④, and ⑤⑥. Fig. 3 suggests that Instant-NVR [9] fails to synthesize time-varying high-frequency wrinkles, and Neural Body [43] synthesizes time-varying yet blurry wrinkles. HumanNeRF [61] synthesizes detailed yet static T-shirt wrinkles, *i.e.*, the wrinkles are almost the same for similar poses, such as ①②, ③④, and ⑤⑥. In contrast, our method renders both high-fidelity and time-varying wrinkles. The comparisons on other subjects S387 and S315 are shown in Fig. 4, which illustrate the effectiveness of our proposed paradigm in dynamics learning.

**Time-varying Appearances with Motion-dependent Shadows on MPII-RDDC.** We compare with baseline methods for novel view synthesis on MPII-RDDC, as shown in Fig. 5. The sequence is captured in a studio with top-down lighting that casts shadows on the human body due to self-occlusions. We notice that the synthesis of shadow can be formulated as learning motion-dependent shadows in the motion representation without the need to




Figure 5. Novel view synthesis of time-varying appearances with both pose and lighting conditioning on MPII-RDDC dataset. The sequence is captured in a studio with top-down lighting that casts shadows on the human performer due to self-occlusion. In Row 1, we specifically focus on synthesizing time-varying shadows (e.g., ① vs. ②, and ③ vs. ④) for different poses with different self-occlusions. In Row 2, we evaluate the synthesis of: 1) time-varying appearances for similar poses occurring in a jump-up-and-down motion sequence, e.g., ⑤ vs. ⑥, 2) shadows ⑦ vs. ⑧, and 3) clothing offsets ⑤ vs. ⑥.




Figure 6. Novel view synthesis of fast motions on AIST++.

explicitly model the lighting. We validate our method from two aspects of the scenario in Fig. 5. 1) In Row 1, Fig. 5 indicates that with the expressive motion representation and physical motion learning, SurMo succeeds in predicting the motion-dependent shadows, such as ①②③④, whereas SOTA HumanNeRF renders almost the same T-shirt appearance. 2) In Row 2, we compare the synthesis of time-varying appearances for similar poses occurring in a jump-up-and-down motion sequence *i.e.*, ⑤ vs. ⑥, which shows that SurMo is capable of predicting the clothes offsets under similar pose with different motion trajectories. In addition, we synthesize the shadows ⑦⑧ in the challenging jump-up-and-down motion sequence. However, HumanNeRF fails to predict the dynamic clothing offsets ⑨.

**Time-varying Appearances for Fast Motions on AIST++.** We also evaluate our methods by rendering humans with fast dance motions (S21 and S13) on AIST++, as shown in Fig. 6, where we compare ours against Neural Body in synthesizing time-varying wrinkles of the T-shirt for different motions. Fig. 6 suggests that Neural Body can only synthesize blurry results partly because Neural Body models pose in a sparse 3D space. Instead, we model motions in a denser body surface space, which enables high-fidelity image synthesis. Note that AIST++ is based

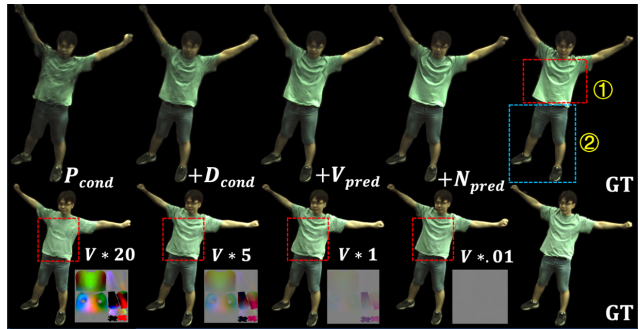


Figure 8. Ablation study of motion conditioning and learning. We focus on the effect of motion conditioning and learning in Row 1, and whether our method learns to decouple the static pose and dynamics (e.g., velocity) from the motion conditioning.

on SMPL tracking with a scaling factor that affects the inverse LBS used in Instant-NVR or HumanNeRF both relying on inverse LBS, and hence we cannot compare them based on the released official code. Quantitative results can be found in the supp. mat.

## 4.2. Ablation Study

**Surface-based Triplane vs. Volumetric Triplane.** We compare the well-adopted volumetric triplane (Vol-Trip) [5] and our proposed surface-based triplane (Surf-Trip) for human modeling on two subjects S313 and S387 from ZJU-MoCap, as shown in Fig. 7. Fig. 7 (left) suggests that the Surf-Trip converges faster in training with a smaller reconstruction loss for both sequences, and the performances are further improved with physical motion learning, e.g., ‘Surf-Trip + ML’. Fig. 7 (right) compares the rendering re-

Convergence of Volumetric Triplane vs. Surface Triplane

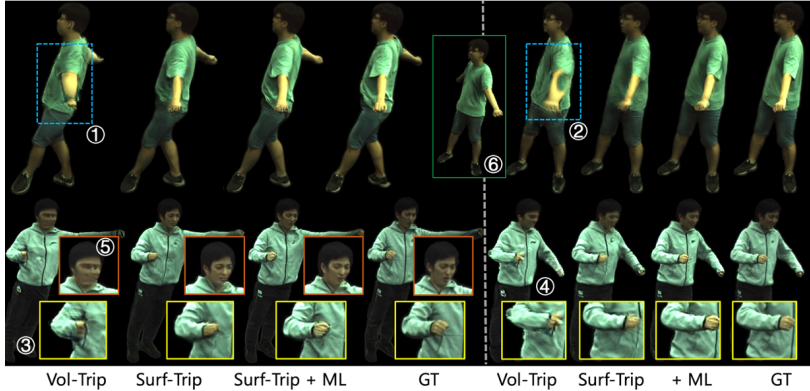
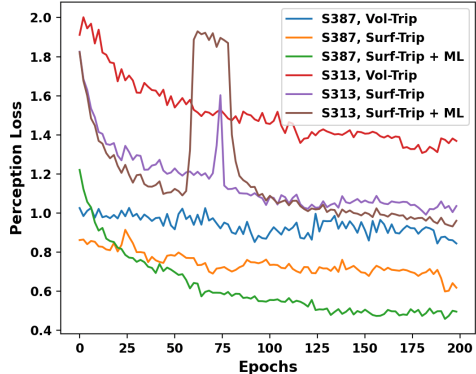


Figure 7. Ablation study of 3D volumetric triplane (Vol-Trip) vs. surface-based triplane (Surf-Trip) for human modeling on subject S313 and S387 from ZJU-MoCap. We focus on the convergence in training (left), and self-occlusion rendering (right). The convergence of S313 is shifted for visualization purposes. Vol-Trip is not effective in handling self-occlusion, *e.g.*, ①②③④ though it performs well in another viewpoint without self-occlusions ⑥. In addition, Vol-Trip cannot synthesize high-quality details for face.

Table 2. Ablation study of motion conditioning and learning.  $P_{cond}$  and  $D_{cond}$  denote the conditioning of pose and dynamics,  $V_{pred}$  and  $N_{pred}$  denote the prediction of surface velocity and surface normal in motion learning.

S313	LPIPS ↓	FID ↓	SSIM ↑	PSNR ↑
$P_{cond}$	.120	104.83	.793	22.781
+ $D_{cond}$	.085	73.674	.834	24.908
+ $V_{pred}$	.069	62.092	.856	25.845
+ $N_{pred}$	.060	50.170	.869	26.654

sults qualitatively, which indicates that Surf-Trip is more effective in handling self-occlusions, whereas Vol-Trip fails ①②③④, though it performs well in another viewpoint without self-occlusions ⑥. In addition, Surf-Trip generates high-fidelity clothing wrinkles and facial details, whereas the face rendering of Vol-Trip is blurry.

**Motion Conditioning and Learning.** We also analyze the effectiveness of our proposed motion conditioning and learning qualitatively and quantitatively. Tab. 2 summarizes the quantitative results, which suggest that the quantitative results are significantly improved by taking as input the dynamics conditioning  $D_{cond}$ . With physical motion learning, *i.e.*, predicting the temporal derivatives surface velocity and spatial derivatives normal at the next timestep, the quantitative results are further improved by a big margin.

The qualitative comparisons are shown in Fig. 8, where in Row 1, we observe higher-fidelity appearances with dynamics conditioning, and physical motion learning for both velocity and normal prediction. In Row 2, we evaluate whether our method learns to decouple the temporal dynamics (*e.g.*, velocity) and static poses from motion conditioning, which corresponds to the rendering of clothing wrinkles of secondary motion (T-shirt ①) and tight parts (*e.g.*, shorts ②). We only change the velocity for each variant

in the ablation study. It suggests that the appearance of the T-shirt varies when the velocity or dynamics are changed, whereas the renderings of tight parts remain the same, such as the head, tight shorts, and shoes. This is consistent with our daily observations. The ablation study illustrates that our method decouples the pose and dynamics, and is capable of generating dynamics-dependent appearances. Note that the appearance changes are small between  $V * 1$  and  $V * 0.01$  when we reduce the velocity, partly because the method is not sensitive to smaller velocity due to the pose estimation errors in the training dataset where consecutive frames may be estimated with the same pose fitting.

## 5. Discussion

We propose SurMo, a new paradigm for learning dynamic humans from videos by jointly modeling temporal motions and human appearances in a unified framework, based on an efficient surface-based triplane. We conduct a systematical analysis of how human appearances are affected by temporal dynamics, and extensive experiments validate the expressiveness of the surface-based triplane in rendering fast motions and motion-dependent shadows. Quantitative experiments illustrate that SurMo achieves the SOTA results on different motion sequences.

## Acknowledgement

This study is supported by the Ministry of Education, Singapore, under its MOE AcRF Tier 2 (MOET2EP20221-0012), NTU NAP, and under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).



## References

- [1] Kfir Aberman, M. Shi, Jing Liao, Dani Lischinski, B. Chen, and D. Cohen-Or. Deep video-based performance cloning. *Computer Graphics Forum*, 38, 2019. 2
- [2] George Borshukov, Dan Piloni, Oystein Larsen, J. P. Lewis, and Christina Tempelaar-Lietz. Universal capture: image-based facial animation for "the matrix reloaded". In *SIGGRAPH '03*, 2003. 1
- [3] Joel Carranza, Christian Theobalt, Marcus A. Magnor, and Hans-Peter Seidel. Free-viewpoint video of human actors. *ACM SIGGRAPH 2003 Papers*, 2003. 1
- [4] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A. Efros. Everybody dance now. *ICCV*, pages 5932–5941, 2019. 2
- [5] Eric Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J. Guibas, Jonathan Tremblay, S. Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3d generative adversarial networks. *ArXiv*, abs/2112.07945, 2021. 2, 4, 5, 7
- [6] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, and Huchuan Lu. Animatable neural radiance fields from monocular rgb video. *ArXiv*, abs/2106.13629, 2021. 2
- [7] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *CVPR*, 2019. 2
- [8] Boyang Deng, John P Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Nasa neural articulated shape approximation. In *ECCV*, 2020. 2
- [9] Chen Geng, Sida Peng, Zhen Xu, Hujun Bao, and Xiaowei Zhou. Learning neural volumetric representations of dynamic humans in minutes. In *CVPR*, 2023. 6
- [10] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [11] A. K. Grigor'ev, Artem Sevastopolsky, Alexander Vakhitov, and Victor S. Lempitsky. Coordinate-based texture inpainting for pose-guided human image generation. *CVPR*, pages 12127–12136, 2019. 2
- [12] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. *ArXiv*, abs/2110.08985, 2021. 2
- [13] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM Transactions on Graphics (TOG)*, 40:1–16, 2021. 2
- [14] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, 2017. 6
- [15] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. *ArXiv*, abs/2112.05637, 2021. 2
- [16] Tao Hu, Geng Lin, Zhizhong Han, and Matthias Zwicker. Learning to generate dense point clouds with textures on multiple categories. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2170–2179, January 2021. 2
- [17] Tao Hu, Kripasindhu Sarkar, Lingjie Liu, Matthias Zwicker, and Christian Theobalt. Egorenderer: Rendering human avatars from egocentric camera images. In *ICCV*, 2021. 2
- [18] Tao Hu, Hongyi Xu, Linjie Luo, Tao Yu, Zerong Zheng, He Zhang, Yebin Liu, and Matthias Zwicker. Hvtr++: Image and pose driven human avatars using hybrid volumetric-textural rendering. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–15, 2023. 2, 3, 6
- [19] T. Hu, Tao Yu, Zerong Zheng, He Zhang, Yebin Liu, and Matthias Zwicker. Hvtr: Hybrid volumetric-textural rendering for human avatars. *3DV*, 2022. 2
- [20] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. *2020 (CVPR)*, pages 3090–3099, 2020. 2
- [21] James T. Kajiya and Brian Von Herzen. Ray tracing volume densities. *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, 1984. 5
- [22] Diederick P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [23] Bernhard Kratzwald, Zhiwu Huang, Danda Pani Paudel, and Luc Van Gool. Towards an understanding of our world by GANing videos in the wild. *arXiv:1711.11453*, 2017. 2
- [24] Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. Learn to dance with aist++: Music conditioned 3d dance generation, 2021. 2
- [25] 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. *TOG*, 40, 2021. 1, 2, 3, 6
- [26] 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. *IEEE Transactions on Visualization and Computer Graphics*, 05 2020. 2
- [27] Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Hyeonwoo Kim, Florian Bernard, Marc Habermann, Wenping Wang, and Christian Theobalt. Neural rendering and reenactment of human actor videos. *ACM Transactions on Graphics (TOG)*, 2019. 2
- [28] M. Loper, Naureen Mahmood, J. Romero, Gerard Pons-Moll, and Michael J. Black. Smpl: a skinned multi-person linear model. *ACM Trans. Graph.*, 34:248:1–16, 2015. 2
- [29] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *NeurIPS*, pages 405–415, 2017. 2
- [30] Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image generation. *CVPR*, 2018. 2
- [31] Qianli Ma, Shunsuke Saito, Jinlong Yang, Siyu Tang, and Michael J. Black. Scale: Modeling clothed humans with a surface codec of articulated local elements. In *CVPR*, 2021. 2
- [32] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J Black. The power of points for modeling humans in clothing. In *ICCV*, 2021. 2
- [33] Lars M. Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy net-

- works: Learning 3d reconstruction in function space. *CVPR*, 2019. 2
- [34] Mateusz Michalkiewicz, Jhony Kaesemodel Pontes, Dominic Jack, Mahsa Baktash, and Anders P. Eriksson. Deep level sets: Implicit surface representations for 3d shape inference. *ArXiv*, 2019.
- [35] Marko Mihajlovic, Yan Zhang, Michael J Black, and Siyu Tang. Leap: Learning articulated occupancy of people. In *CVPR*, 2021. 2
- [36] Natalia Neverova, Riza Alp Güler, and Iasonas Kokkinos. Dense pose transfer. *ECCV*, 2018. 2
- [37] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. *CVPR*, pages 11448–11459, 2021. 2
- [38] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *IEEE/CVF ICCV*, 2021. 2
- [39] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. Stylesdf: High-resolution 3d-consistent image and geometry generation. *ArXiv*, abs/2112.11427, 2021. 2
- [40] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. Npms: Neural parametric models for 3d deformable shapes. In *IEEE/CVF ICCV*, 2021. 2
- [41] Jeong Joon Park, Peter R. Florence, Julian Straub, Richard A. Newcombe, and S. Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. *2019 (CVPR)*, pages 165–174, 2019. 2
- [42] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, 2021. 2, 3
- [43] 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, 3, 6
- [44] Sergey Prokudin, Michael J. Black, and Javier Romero. Smpix: Neural avatars from 3d human models. *WACV*, 2021. 2
- [45] Albert Pumarola, Antonio Agudo, Alberto Sanfeliu, and Francesc Moreno-Noguer. Unsupervised person image synthesis in arbitrary poses. In *CVPR*, June 2018. 2
- [46] Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. Anr: Articulated neural rendering for virtual avatars. *CVPR*, pages 3721–3730, 2021. 2
- [47] Edoardo Remelli, Timur M. Bagautdinov, Shunsuke Saito, Chenglei Wu, Tomas Simon, Shih-En Wei, Kaiwen Guo, Zhe Cao, Fabián Prada, Jason M. Saragih, and Yaser Sheikh. Drivable volumetric avatars using texel-aligned features. *ACM SIGGRAPH*, 2022. 2, 3, 6
- [48] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. *IEEE/CVF ICCV*, pages 2304–2314, 2019. 2
- [49] Shunsuke Saito, Tomas Simon, Jason M. Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. *2020 (CVPR)*, pages 81–90, 2020.
- [50] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. *2021 (CVPR)*, pages 2885–2896, 2021. 2
- [51] Kripasindhu Sarkar, Dushyant Mehta, Weipeng Xu, Vladislav Golyanik, and Christian Theobalt. Neural re-rendering of humans from a single image. In *ECCV*, 2020. 2, 6
- [52] Aliaksandr Siarohin, Enver Sangineto, Stephane Lathuiliere, and Nicu Sebe. Deformable GANs for pose-based human image generation. In *CVPR*, 2018. 2
- [53] Shih-Yang Su, Frank Yu, Michael Zollhoefer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. In *NeurIPS*, 2021. 2
- [54] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason M. Saragih, Matthias Nießner, Rohit Pandey, S. Fanello, Gordon Wetzstein, Jun-Yan Zhu, Christian Theobalt, Maneesh Agrawala, Eli Shechtman, Dan B. Goldman, and Michael Zollhofer. State of the art on neural rendering. *Computer Graphics Forum*, 2020. 1
- [55] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38, 2019. 2
- [56] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-gif: Neural generalized implicit functions for animating people in clothing. In *ICCV*, 2021. 2
- [57] Shaofei Wang, Marko Mihajlovic, Qianli Ma, Andreas Geiger, and Siyu Tang. Metaavatar: Learning animatable clothed human models from few depth images. *NeurIPS*, 2021. 2
- [58] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdfs. In *European Conference on Computer Vision*, 2022. 6
- [59] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *NeurIPS*, 2018. 2
- [60] Zhou Wang, A. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13:600–612, 2004. 5, 6
- [61] Chung-Yi Weng, Brian Curless, Pratul P. Srinivasan, Jonathan T. Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. *ArXiv*, abs/2201.04127, 2022. 1, 2, 3, 6
- [62] 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. *ACM SIGGRAPH*, 2011. 1
- [63] Jae Shin Yoon, Duygu Ceylan, Tuanfeng Y. Wang, Jingwan Lu, Jimei Yang, Zhixin Shu, and Hyunjung Park. Learning motion-dependent appearance for high-fidelity rendering of dynamic humans from a single camera. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3397–3407, 2022. 2, 3
- [64] Richard Zhang, Phillip Isola, Alexei A. Efros, E. Shechtman,

and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. *CVPR*, pages 586–595, 2018. 6

- [65] Yang Zheng, Ruizhi Shao, Yuxiang Zhang, Tao Yu, Zerong Zheng, Qionghai Dai, and Yebin Liu. Deepmulticap: Performance capture of multiple characters using sparse multiview cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6239–6249, 2021. 2
- [66] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *TPAMI*, PP, 2021. 2
- [67] Peng Zhou, Lingxi Xie, Bingbing Ni, and Qi Tian. Cips-3d: A 3d-aware generator of gans based on conditionally-independent pixel synthesis. *ArXiv*, 2021. 2
- [68] Zhen Zhu, Tengeng Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. Progressive pose attention transfer for person image generation. In *CVPR*, pages 2347–2356, 2019. 2