

HOSNeRF: Dynamic Human-Object-Scene Neural Radiance Fields from a Single Video

Jia-Wei Liu^{1*}, Yan-Pei Cao^{2†}, Tianyuan Yang¹, Zhongcong Xu¹, Jussi Keppo^{4,5},
 Ying Shan², Xiaohu Qie³, Mike Zheng Shou^{1†}

¹ Show Lab, National University of Singapore ² ARC Lab, ³ Tencent PCG

⁴ Business School, ⁵ Institute of Operations Research and Analytics, National University of Singapore

Abstract

We introduce *HOSNeRF*, a novel 360° free-viewpoint rendering method that reconstructs neural radiance fields for dynamic human-object-scene from a single monocular in-the-wild video. Our method enables pausing the video at any frame and rendering all scene details (dynamic humans, objects, and backgrounds) from arbitrary viewpoints. The first challenge in this task is the complex object motions in human-object interactions, which we tackle by introducing the new object bones into the conventional human skeleton hierarchy to effectively estimate large object deformations in our dynamic human-object model. The second challenge is that humans interact with different objects at different times, for which we introduce two new learnable object state embeddings that can be used as conditions for learning our human-object representation and scene representation, respectively. Extensive experiments show that *HOSNeRF* significantly outperforms SOTA approaches on two challenging datasets by a large margin of 40% ~ 50% in terms of LPIPS. The code, data, and compelling examples of 360° free-viewpoint renderings from single videos: <https://showlab.github.io/HOSNeRF>.

1. Introduction

Video reconstruction and free-viewpoint rendering offer innovative opportunities for creating immersive experiences, encompassing virtual reality, telepresence, metaverse, and 3D animation production. While reconstructing videos has the potential to enhance user engagement and provide more realistic environments, it also poses significant challenges in terms of monocular viewpoints [55, 48] and complicated human-environment interactions [50, 14].

In recent years, remarkable progress has been made in

novel view synthesis, particularly since the introduction of Neural Radiance Fields (NeRF) [27]. While initially limited to reconstructing static 3D scenes based on multi-view images, subsequent studies have proposed various approaches to address the challenge of dynamic view synthesis [34, 29, 30, 47, 20, 8]. NeRF-based techniques have evolved to either incorporate deformation fields that map dynamic fields to canonical NeRF spaces [34, 29, 30, 47], or model dynamic scenes as 4D spatio-temporal radiance fields [20, 8]. While these approaches have shown promising results in dynamic view synthesis, they are limited to simple deformations. Another line of research is specifically designed for dynamic neural human modeling that relies on estimated human poses as a priori information [32, 50]. Recently, Neuman [14] combines pose-driven dynamic human models with static scene models for representing dynamic human-centric scenes.

However, none of the aforementioned techniques can accurately reconstruct challenging monocular videos with fast and complex human-object-scene motions and interactions, as shown in Fig. 1(d). This is due to two particular challenges listed below. To tackle them, we introduce a novel method called **Human-Object-Scene Neural Radiance Fields (HOSNeRF)**.

i) Complex object motions in human-object interactions. In contrast to the simple motions that can be modeled by general deformation modules [29, 30], the object motion during human-object interaction is often drastic and composed of various atomic motions (*e.g.*, play tennis). To tackle this challenge, we propose new object bones that are attached to the human skeleton hierarchy to estimate human-object deformations in a coarse-to-fine manner for our dynamic human-object model. The object bones and underlying object linear blend skinning (object LBS) allow for the accurate estimation of objects' deformations through the relative transformations in the kinematic tree of the skeleton hierarchy.

*Work is partially done during internship at ARC Lab, Tencent PCG.

†Corresponding Authors.

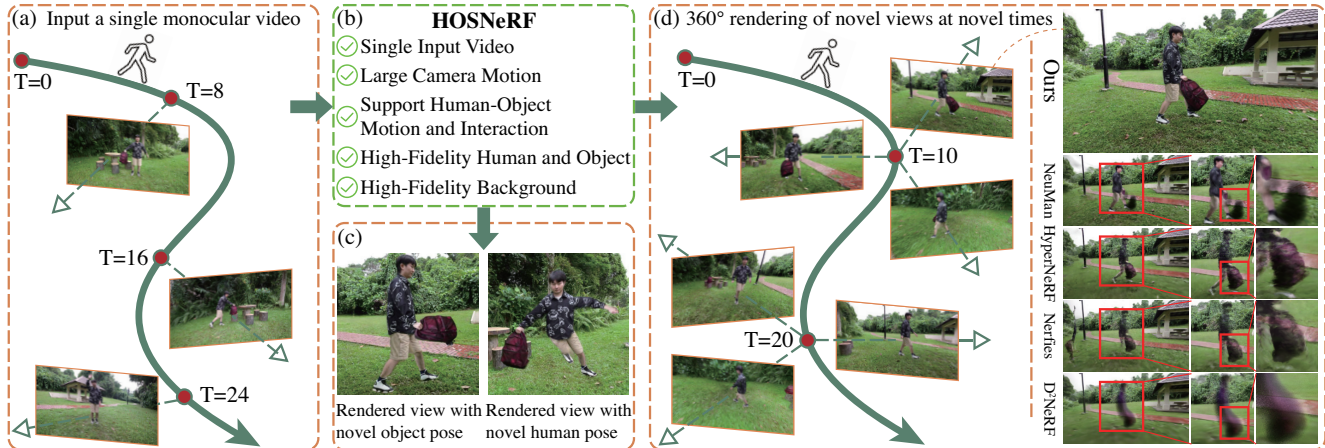


Figure 1: Our HOSNeRF (b) takes a single monocular in-the-wild video (a) as input, and creates high-fidelity 360° free-viewpoint rendering of all scene details (dynamic human body, objects, and background) at any time (d). Our method enables rendering views with novel object poses and novel human poses as shown in (c), and produces high-fidelity dynamic novel view synthesis results at novel timesteps, with significant improvements over SOTA approaches as shown in (d).

ii) Humans interact with different objects at different times. The above human-object model is designed for humans interacting with the same object over time. But when the person puts down the current object or picks a new object, it is not clear how to dynamically remove/add such objects in the static background model and the human-object model whose canonical space is static. To handle this challenge, we introduce two new learnable object state embeddings that can be used as conditions for learning our human-object representation and scene representation, respectively.

Finally, we systematically explore and identify effective training objectives and strategies for our proposed HOSNeRF, including deformation cycle consistency, optical flow supervisions, and foreground-background rendering. On two challenging datasets collected by ourselves and NeuMan [14], our HOSNeRF achieves high-fidelity dynamic novel view synthesis results and enables pausing the monocular video at any time and rendering all scene details (dynamic humans, objects, and backgrounds) from arbitrary viewpoints, as shown in Fig. 1(d).

In summary, our main contributions are:

- We present a novel framework of HOSNeRF, the first work to achieve 360° free-viewpoint high-fidelity novel view synthesis for dynamic scenes with human-environment interactions from a single video.
- We propose the object bones and state-conditional representations to handle the non-rigid motions and interactions of humans, objects, and the environment.
- Extensive experiments show that HOSNeRF significantly outperforms SOTA approaches on two challenging datasets by 40% ~ 50% in terms of LPIPS.

2. Related Work

2.1. Dynamic Human Modeling

Dynamic human modeling has shown promising results in utilizing various representations such as point clouds [1, 51], meshes [10, 24], voxels [25, 39], and neural implicit functions [12, 35], with models like SMPL [26, 33] being commonly used for parameterizing the human body. Since the introduction of NeRF [27], neural human representation [32, 36, 42, 54, 23] has achieved remarkable progress on representing dynamic human bodies from sparse-view videos. Among them, Neural Actor [23] and Neural Body [32] pioneer in combining NeRF [27] with SMPL deformable meshes to represent human bodies with complex motions. Subsequent works have further improved on the generalizability [17, 4, 9] and animatability [31, 18] of human bodies. To support multi-person modeling, recent works [55, 37] have proposed to segment each human into 3D bounding boxes and learn a separate layered dynamic NeRF for each person. Other works [46, 43, 15, 44] are specifically designed to reconstruct the dynamic human and object with RGB-D or multi-view videos as inputs. They track the human and object pose, and separately reconstruct them with volumetric fusion [43], neural texture blending [44], or neural rendering [15, 56].

Despite achieving promising results, these approaches require multi-view videos or RGB-D as input, limiting their real-world applications. To solve this problem, HumanNeRF [50] is proposed to represent moving humans from a monocular video by the human pose-driven deformation module and canonical space. NeuMan [14] is the first successful attempt at reconstructing both the dynamic human and static background from a single video. However, NeuMan [14] does not support human-environment interactions

and performs poorly at large camera motions.

2.2. Dynamic View Synthesis for General Scenes

Most prior approaches on dynamic scene modeling require synchronized multi-view videos [58, 19, 41, 5, 55, 48] or depth [59, 13, 28] as input. Recent studies have built upon NeRF [27] to reconstruct dynamic neural radiance fields from monocular videos by either learning a deformation field that maps dynamic observation to canonical field [34, 29, 30, 47] or building 4D spatio-temporal radiance fields [53, 20, 8]. Among them, Nerfies [29] associates latent codes with the deformation field and HyperNeRF [30] represents motion in a high-dimensional space. D²NeRF [52] builds upon HyperNeRF [30] and further decouples the dynamic components from the static background, and represents them separately with a HyperNeRF [30] and NeRF [27]. DynIBaR [21] proposes a motion-adjusted multi-view feature aggregation module to synthesize new viewpoints by aggregating features from nearby views. Other studies have introduced voxel grids [6, 22, 40] or planar representations [7, 3] for fast dynamic radiance fields reconstruction. While these approaches have achieved high-fidelity dynamic view synthesis results, they are restricted to simple scene deformations. In contrast, our HOSNeRF is capable of representing significant human-object motions and interactions in complex environments.

3. Method

3.1. Preliminaries

HumanNeRF [50] has been recently introduced to represent a moving person with a NeRF [27]-based canonical space Ψ_c that maps 3D points to color \mathbf{c} and density d , and a human pose-guided deformation field Ψ_d that transforms deformed points \mathbf{x}_d^i from the deformed space to canonical points \mathbf{x}_c^i in the canonical space (i omitted for simplicity).

$$\Psi_c(\gamma(\mathbf{x}_c)) \mapsto (\mathbf{c}, \sigma), \quad \Psi_d(\mathbf{x}_d, \mathcal{J}, \mathcal{R}) \mapsto (\mathbf{x}_c), \quad (1)$$

where $\gamma(\mathbf{x})$ is the standard positional encoding function, and $\mathcal{J} = \{\mathbf{J}_i\}$ and $\mathcal{R} = \{\omega_i\}$ represent the 3D human joints and the local joint axis-angle rotations.

The deformation field Ψ_d is decomposed into the coarse human skeleton-driven deformation Ψ_d^{coarse} , and the fine non-rigid deformation conditioned on human poses Ψ_d^{fine} :

$$\mathbf{x}_c' = \Psi_d^{\text{coarse}}(\mathbf{x}_d, \mathcal{J}, \mathcal{R}), \quad \mathbf{x}_c = \mathbf{x}_c' + \Psi_d^{\text{fine}}(\mathbf{x}_c', \mathcal{R}). \quad (2)$$

Mip-NeRF 360 [2] is designed to synthesize realistic views for highly intricate, unbounded real-world static scenes. To render pixel colors, the casted rays are split into a set of intervals $T_i = [t_i, t_{i+1})$. The mean and covariance of the conical frustum corresponding to every interval are computed as $(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \mathbf{r}(T_i)$. To parameterize the Gaussian

parameters for unbounded scenes, Mip-NeRF 360 [2] further proposes a contraction function $f(\mathbf{x})$ that distributes distant points proportionally to disparity,

$$f(\mathbf{x}) = \begin{cases} \mathbf{x} & \|\mathbf{x}\| \leq 1 \\ \left(2 - \frac{1}{\|\mathbf{x}\|}\right) \left(\frac{\mathbf{x}}{\|\mathbf{x}\|}\right) & \|\mathbf{x}\| > 1 \end{cases}. \quad (3)$$

Then, $f(\mathbf{x})$ is applied to $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ as follows:

$$(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}) = \left(f(\boldsymbol{\mu}), \mathbf{J}_f(\boldsymbol{\mu}) \boldsymbol{\Sigma} \mathbf{J}_f(\boldsymbol{\mu})^T\right), \quad (4)$$

where $\mathbf{J}_f(\boldsymbol{\mu})$ is the Jacobian of f at $\boldsymbol{\mu}$. The contracted Gaussian parameters $(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})$ are further encoded through the integrated positional encoding (IPE) [2]:

$$\hat{\gamma}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}) = \left\{ \left[\begin{array}{c} \sin(2^\ell \hat{\boldsymbol{\mu}}) \exp\left(-2^{2\ell-1} \text{diag}\left(\hat{\boldsymbol{\Sigma}}\right)\right) \\ \cos(2^\ell \hat{\boldsymbol{\mu}}) \exp\left(-2^{2\ell-1} \text{diag}\left(\hat{\boldsymbol{\Sigma}}\right)\right) \end{array} \right] \right\}_{\ell=0}^{L-1}, \quad (5)$$

and the color and density of intervals can be obtained as

$$\Psi_s\left(\hat{\gamma}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})\right) \mapsto (\mathbf{c}, \sigma), \quad (6)$$

where Ψ_s is the scene NeRF MLP [2].

3.2. Dynamic Human-Object Model

Motivation. With estimated 3D human poses a priori, HumanNeRF [50] is effective at modeling moving people, but can only encode body parts, lacking the capability to model additional structures (*e.g.*, objects held by the person) and hence not suitable for challenging scenarios containing complex human-object interactions.

A naive approach would be attaching the objects directly to the interacting body parts (*e.g.*, hands) so that the object transformations can be queried via the skeletal motion of the corresponding human joints. However, this approach cannot represent the relative transformations between the object and its interacting body part, making it difficult to model large objects, such as the suitcase shown in Fig. 6.

Object Bones. To address this limitation, based on the conventional human skeleton hierarchy, we introduce a new ‘‘object bone’’ for each hand, where the starting point of the object bone is connected to the corresponding hand joint. As shown in Fig. 2 (a), two new object bones (green color) are connected to the hand joints, which represent the relative rotation and translation with respect to their parent hand joints. In potential, the number, bone size, and attaching joints of object bones can also be flexibly customized based on different human-object interaction scenarios. Therefore, our 3D human-object pose consists of 3D joints $\mathcal{J} = \{\mathbf{J}_{\text{human}}, \mathbf{J}_{\text{object}}\}$ and their axis-angle rotations $\mathcal{R} = \{\omega_{\text{human}}, \omega_{\text{object}}\}$.

Object LBS. To drive the articulated motions of the objects and human body, we learn a linear blend skinning

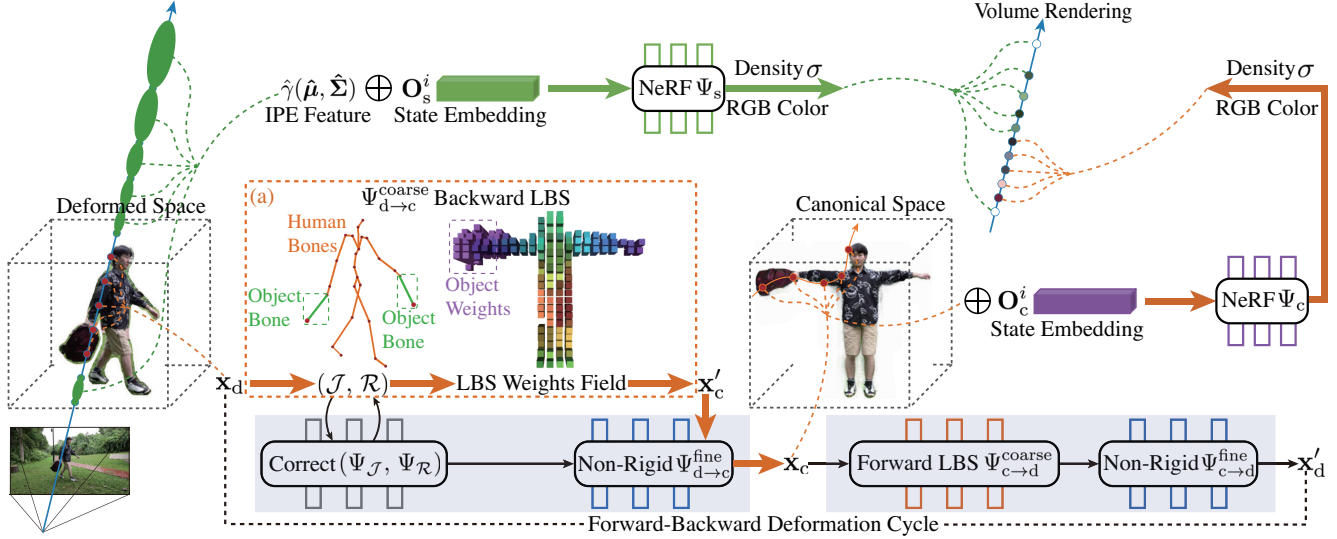


Figure 2: **Overview of our method.** (1) **Orange flowchart:** The deformation from deformed points to canonical points are effectively estimated by the human-object backward LBS (a) and non-rigid deformation module, and their properties (i.e., density, color) can be obtained by querying the state-conditional canonical space. (2) **Green flowchart:** The background intervals sampled from a deformed frame are concatenated with object state embeddings for querying properties through the state-conditional scene model. (3) Novel views can be accordingly synthesized by volume rendering for re-ordered properties.

(LBS) weights field during training, as shown in Fig. 2 (a), which encodes the influence regions of object bones and human bones. Following HumanNeRF [50], we pack the human and object bone blend weights and additional background weight into a single volume with $K + 1$ channels, i.e., $\mathcal{W}_{c \rightarrow d}(\mathbf{x}_c) = \{w_{c \rightarrow d}^i(\mathbf{x}_c)\}$, and generate the volume from a learnable latent code using a CNN.

For instance, in Fig. 2, the subject holds a backpack in the right hand, so the learned LBS weight field only encodes the influence region of the right object bone with a similar shape to the backpack, while the left object bone has no influence on skeleton motion. Therefore, this learnable LBS weights field enables our object bones to be generalizable to various human-object interaction scenarios with left, right, or both hands. In addition, our human-object skeleton hierarchy enables HOSNeRF to support rendering with novel object pose and novel human pose, as illustrated in Fig. 1(c).

Forward-Backward Deformation Framework. To improve the smoothness and consistency of human-object deformations, we propose to leverage the *cycle consistency* between forward and backward human-object deformations. The backward deformation first transforms sampled deformed points \mathbf{x}_d to their corresponding canonical points \mathbf{x}_c , which are further mapped back to the deformed space $\hat{\mathbf{x}}_d$ through the forward deformation module.

As illustrated in Fig. 2(a), with the proposed human-object skeleton definition and LBS weight fields, the coarse backward skeleton deformation $\Psi_{d \rightarrow c}^{\text{coarse}}$ is defined similar to

HumanNeRF [50],

$$\Psi_{d \rightarrow c}^{\text{coarse}}(\mathbf{x}_d, \mathcal{J}, \mathcal{R}) = \sum_{i=1}^K w_{d \rightarrow c}^i(\mathbf{x}_d) (\mathbf{R}_{d \rightarrow c}^i \mathbf{x}_d + \mathbf{t}_{d \rightarrow c}^i), \quad (7)$$

where K is the number of human-object bones, $w_{d \rightarrow c}^i$ is the blend weight for the i -th bone at the deformed space, and $\mathbf{R}_{d \rightarrow c}^i$ and $\mathbf{t}_{d \rightarrow c}^i$ are the backward rotation and translation that are explicitly derived from the human-object pose.

To model the non-rigid motion that can not be represented by skeleton motion, we additionally define the backward fine non-rigid motion $\Psi_{d \rightarrow c}^{\text{fine}}$ conditioned on the human-object poses [50]:

$$\Delta \mathbf{x}_{d \rightarrow c} = \Psi_{d \rightarrow c}^{\text{fine}}(\gamma(\Psi_{d \rightarrow c}^{\text{coarse}}(\mathbf{x}_d, \mathcal{J}, \mathcal{R})), \mathcal{R}). \quad (8)$$

Therefore, \mathbf{x}_d can be mapped to the canonical space by

$$\mathbf{x}_c = \Psi_{d \rightarrow c}^{\text{coarse}}(\mathbf{x}_d, \mathcal{J}, \mathcal{R}) + \Delta \mathbf{x}_{d \rightarrow c}. \quad (9)$$

Accordingly, the forward coarse skeleton motion driven by the human-object pose can be defined as

$$\Psi_{c \rightarrow d}^{\text{coarse}}(\mathbf{x}_c, \mathcal{J}, \mathcal{R}) = \sum_{i=1}^K w_{c \rightarrow d}^i(\mathbf{x}_c) (\mathbf{R}_{c \rightarrow d}^i \mathbf{x}_c + \mathbf{t}_{c \rightarrow d}^i), \quad (10)$$

where $w_{c \rightarrow d}^i$ is the canonical blend weight for the i -th bone, $\mathbf{R}_{c \rightarrow d}^i$ and $\mathbf{t}_{c \rightarrow d}^i$ are the forward rotation and translation obtained from the human-object pose. In addition, the forward fine non-rigid motion $\Psi_{c \rightarrow d}^{\text{fine}}$ is

$$\Delta \mathbf{x}_{c \rightarrow d} = \Psi_{c \rightarrow d}^{\text{fine}}(\gamma(\Psi_{c \rightarrow d}^{\text{coarse}}(\mathbf{x}_c, \mathcal{J}, \mathcal{R})), \mathcal{R}). \quad (11)$$

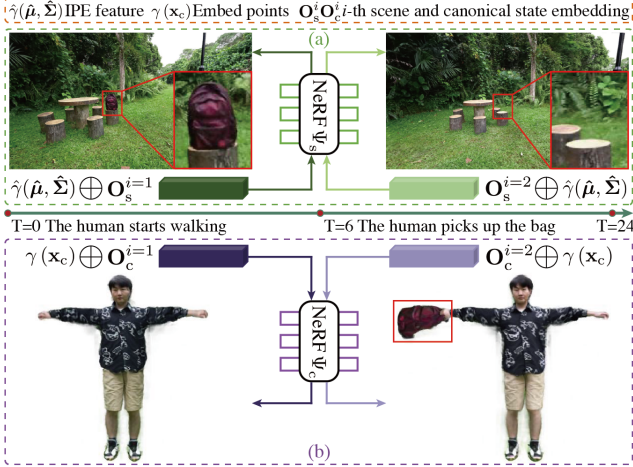


Figure 3: State-conditional representation for the scene model (a) and human-object canonical space (b).

Therefore, the canonical points \mathbf{x}_c can be transformed back to the deformed scene by

$$\hat{\mathbf{x}}_d = \Psi_{c \rightarrow d}^{\text{coarse}}(\mathbf{x}_c, \mathcal{J}, \mathcal{R}) + \Delta \mathbf{x}_{c \rightarrow d}. \quad (12)$$

3.3. State-Conditional Representation

Motivation. In complex dynamic scenes, humans can interact with different objects at different timesteps. As the example in Fig. 3 shows, the human picks up the bag at $T = 6$ and puts it down at $T = 24$, resulting in 3 object states: the bag in the background during $T \in [0, 6]$, the bag held by the human during $T \in [6, 24]$, the bag at a new position in the background when $T > 24$. Such object state changes prohibit us from directly using the static background model (*i.e.*, Mip-NeRF 360 [2]) and the dynamic human-object model proposed in Sec. 3.2.

A naive solution is to train a separate network for each object state. However, this solution is limited because each network is trained only on a short segment (*e.g.*, $T \in [0, 6]$) instead of the whole video. As a result, some regions often are difficult to reconstruct because they are not observed in the short segment. But these regions could have already been captured by video segments of other object states.

To make full use of all segments to jointly train one model shared for different states, we introduce learnable state embeddings to represent different object states. As depicted in Fig. 3, we have two state embeddings respectively for our scene model and human-object model, which are explained in detail as follows.

State-Conditional Scene Model. As shown in Fig. 3(a), we develop a state-conditional Mip-NeRF 360 [2] to represent the static scene with temporal object transitions. In a dynamic scene with N object states, we define N learnable state embeddings $\mathcal{O}_s = \{\mathbf{O}_s^i\}$ ($i = 1, 2, \dots, N$) to represent states information. Therefore, at state i , we concatenate the IPE features $\hat{\gamma}(\hat{\mu}, \hat{\Sigma})$ (Eq. (5)) of ray intervals with

the state embedding \mathbf{O}_s^i as input to the scene MLP Ψ_s for querying the scene color and density.

$$\Psi_s \left(\text{concat} \left(\hat{\gamma}(\hat{\mu}, \hat{\Sigma}), \mathbf{O}_s^i \right) \right) \mapsto (\mathbf{c}, \sigma). \quad (13)$$

This representation enables a shared scene model across the video with multiple object state transitions.

State-Conditional Dynamic Human-Object Model. In Sec. 3.2, we introduce the human-object pose-driven deformation module and the human-object canonical space for modeling dynamic human-objects. However, the reconstructed canonical space is insufficient in representing the temporal changes in object geometry and appearance when the subject interacts with new objects. To address this limitation, we condition the canonical space on the object states, as shown in Fig. 3(b). In particular, we employ an object state-conditional canonical space and optimize a shared blend weights field across all states.

In a dynamic scene with N object states, we define N learnable state embeddings $\mathcal{O}_c = \{\mathbf{O}_c^i\}$ ($i = 1, 2, \dots, N$) to represent object states in the canonical space. Therefore, at object state i , we concatenate the positionally encoded canonical points in Eq. (1) with the state embedding \mathbf{O}_c^i as input to the canonical MLP for querying the human-object color and density.

$$\Psi_c \left(\text{concat} \left(\gamma(\mathbf{x}_c), \mathbf{O}_c^i \right) \right) \mapsto (\mathbf{c}, \sigma). \quad (14)$$

3.4. Training

Rendering. To render the pixel color, we shoot two rays ($\mathbf{r}_s, \mathbf{r}_{ho}$) and sample intervals (*i.e.*, ray points) to query the scene model (using \mathbf{r}_s) and dynamic human-object model (with \mathbf{r}_{ho}) respectively in their coordinate systems. After aligning the coordinates (see supplementary materials), we transform the sampled 3D points and their queried properties from the dynamic human-object space to the scene space, and re-order all sampled properties based on their distance from the camera center, as shown in Fig. 2. Therefore, the color of a pixel can be calculated through volume rendering, *i.e.*, by integrating the re-ordered properties along the ray [27]:

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - e^{-\sigma_i \delta_i}) \mathbf{c}_i, T_i = e^{-\sum_{j=1}^{i-1} \sigma_j \delta_j}. \quad (15)$$

Directly fusing \mathbf{r}_s and \mathbf{r}_{ho} may yield minor artifacts on overlapped ground regions due to intersections of human-object points and scene points on short-length ground rays. To solve this issue, we estimate the foreground mask during training to separate the background rays from foreground ones. This allows us to focus on combining \mathbf{r}_s and \mathbf{r}_{ho} at the overlapped foreground regions, thereby reducing artifacts.

Training Objectives. The training of HOSNeRF consists of three stages. In the first stage, we mask out the dynamic human-object regions and train the state-conditional

Scene	No. of objects	No. of states
BACKPACK	1	3
TENNIS	2	3
SUITCASE	1	4
PLAYGROUND	3	5
DANCE	3	7
LOUNGE	3	5

Table 1: Details of our collected dataset.

scene model. Then, we train the state-conditional dynamic human-object model on the dynamic human-object regions. In the third stage, we combine these two models and further finetune the complete HOSNeRF for all image pixels.

Given the single video with calibrated poses, the first stage of HOSNeRF is optimized by minimizing the photometric MSE loss and the regularization losses proposed by Mip-NeRF 360 [2] to avoid background collapse. In the second and third stages of HOSNeRF, we utilize the photometric MSE loss, patched-based perceptual LPIPS [57] loss, forward-backward deformation cycle consistency, and indirect optical flow supervisions.

The deformation cycle consistency is enforced by minimizing the distance between estimated deformed points of Eq. (12) and sampled deformed points \mathbf{x}_d ,

$$\mathcal{L}_{\text{Cycle}} = \frac{1}{2N} \sum_{i=1}^N \|\mathbf{x}_d^i - \hat{\mathbf{x}}_d^i\|_2^2. \quad (16)$$

We also incorporate 2D optical flow as indirect supervision by minimizing the error between the induced flow and the estimated flow (details in the supplementary material). Therefore, the overall training objective of HOSNeRF is:

$$\begin{aligned} \mathcal{L} = & \omega_{\text{MSE}} \cdot \mathcal{L}_{\text{MSE}} + \omega_{\text{LPIPS}} \mathcal{L}_{\text{LPIPS}} \\ & + \omega_{\text{Cycle}} \cdot \mathcal{L}_{\text{Cycle}} + \omega_{\text{Flow}} \cdot \mathcal{L}_{\text{Flow}}, \end{aligned} \quad (17)$$

where ω_{MSE} , ω_{LPIPS} , ω_{Cycle} , ω_{Flow} are loss weights.

4. Experiments

4.1. Dataset

To benchmark the reconstruction of monocular videos with dynamic human-object-scenes, we collect a new dataset with 6 scenes[†]. The collected dataset features various types of human-object-scene interactions in indoor and outdoor scenarios, with up to 3 interacted objects and 7 states for a single video, as reported in Tab. 1. The duration of collected videos varies from 60s to 120s, and we extract [300, 400] frames for each video, where 16 frames at equal intervals are selected as novel views at novel timesteps as the test set and the remaining frames are the train set.

[†]Our data collection process has been approved by the institutional review board and participants have given consents for public data release.

We extensively evaluate the HOSNeRF on our collected dataset and NeuMan dataset [14] that consists of 6 short human walking sequences. Since NeuMan dataset does not involve human-object interactions, HOSNeRF can be flexibly customized to HOSNeRF by removing object bones and setting the object state to 1 when evaluated on NeuMan dataset.

4.2. Implementation Details

Human interacts with various objects using hands in our dataset, we therefore introduce 2 new object bones to the left and right hand joints of the 3D human skeleton for all scenes. As a result, there are in total $K = 26$ bones in our human-object skeleton. In order to obtain the human-object pose prior, we first utilize a pretrained ROMP [45] model to estimate the 3D human pose, and use Mask-RCNN [11] to estimate human and objects masks. However, since it is non-trivial to estimate the object pose for in-the-wild videos, we initialize the left and right object bones as the standard extended bones with no relative rotations with respect to human hands, and initialize their length by referencing the length of the object relative to the arm. During training, we further refine the human-object poses (\mathcal{J} , \mathcal{R}) with a MLP-based pose correction module ($\Psi_{\mathcal{J}}$, $\Psi_{\mathcal{R}}$) [50]. The overall three-stage training of HOSNeRF takes about 5 days on 4 Tesla V100 GPUs. Please see the supplementary material for the comparison of training time with other approaches and discussions about estimating object masks.

4.3. Experimental Results

Baselines and Evaluation Metrics. To demonstrate the performance of HOSNeRF, we make comparisons with various types of SOTA approaches, including (1) NeuMan [14], a human pose-based method; (2) HyperNeRF [30] and (3) Nerfies [29], dynamic radiance fields for general scenes; (4) D²NeRF [52], a static-dynamic decomposition method; (5) K-Planes [7], an explicit radiance field method. For quantitative comparison, peak signal-to-noise ratio (PSNR), structural similarity index (SSIM) [49], and Learned Perceptual Image Patch Similarity (LPIPS) [57] are employed as evaluation metrics.

Evaluation on HOSNeRF Dataset. We report the quantitative comparisons (PSNR, SSIM, and LPIPS with VGG [38] backbone) of 6 scenes for all approaches in Tab. 2. For all SOTA approaches, we utilize their highest configurations for a fair comparison. As shown in Tab. 2, our HOSNeRF achieves the best performance in terms of all metrics, except the PSNR for the Lounge scene (a metric known to favor smooth/blurry results [57]). The improvement of HOSNeRF is particularly significant in terms of the LPIPS, with an average of nearly 40% gain over SOTA approaches. Fig. 4 visualizes the qualitative comparison of HOSNeRF over SOTA approaches on novel views at novel timesteps, where HOSNeRF achieves substantially better visual qual-




Figure 4: Qualitative comparisons of HOSNeRF and SOTA approaches on HOSNeRF dataset.

ity than other approaches for all scenes. HOSNeRF is able to produce high-fidelity details close to ground truths for all scene contents, i.e., dynamic human bodies, objects, and backgrounds. In contrast, existing approaches tend to synthesize much blurrier images with missing components. Please see the supplementary video for more results on 360° free-viewpoint bullet-time videos rendered by HOSNeRF.

Evaluation on NeuMan Dataset [14]. We report quantitative metrics (PSNR, SSIM, and LPIPS with AlexNet [16] backbone) over 6 scenes of the NeuMan dataset in Tab. 3, where the metrics for NSFF [20], HyperNeRF [30], and NeuMan [14] are borrowed from NeuMan [14]. As shown

in Tab. 3, our HOSNeRF achieves the best performance in terms of all metrics except the PSNR for the Citron scene, and significantly improves SOTA approaches by a large margin of 50% in terms of LPIPS. This further demonstrates the effectiveness and flexibility of HOSNeRF to model various types of dynamic human-(objects)-scenes. Fig. 5 visualizes the qualitative comparisons between HOSNeRF and NeuMan [14] for novel view synthesis at novel timesteps, where HOSNeRF achieves better visual quality over NeuMan [14] for the rendered human bodies and backgrounds in all scenes due to the superiority of our human model and background model. Our HOSNeRF also achieves bet-

	BACKPACK			TENNIS			SUITCASE			PLAYGROUND			DANCE			LOUNGE		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
K-Planes [7]	19.05	0.345	0.557	19.31	0.690	0.536	18.64	0.575	0.602	17.92	0.446	0.635	18.17	0.328	0.623	24.21	0.824	0.453
D ² NeRF [52]	20.52	0.384	0.608	23.97	0.737	0.540	20.99	0.592	0.645	21.23	0.510	0.616	19.92	0.407	0.647	27.13	0.858	0.509
Nerfies [29]	19.56	0.515	0.559	22.12	0.818	0.443	19.01	0.643	0.555	21.14	0.726	0.533	19.37	0.674	0.524	25.90	0.914	0.342
HyperNeRF [30]	19.62	0.359	0.587	21.26	0.742	0.510	19.41	0.604	0.607	21.67	0.558	0.578	19.30	0.412	0.601	27.25	0.925	0.332
NeuMan [14]	21.21	0.479	0.478	23.17	0.768	0.442	20.84	0.611	0.551	21.46	0.546	0.551	21.19	0.529	0.490	28.40	0.917	0.341
Ours (base)	21.51	0.764	0.271	24.02	0.910	0.326	21.10	0.829	0.395	22.20	0.787	0.348	21.84	0.785	0.266	26.54	0.962	0.243
Ours w/ state	22.38	0.786	0.252	23.98	0.910	0.323	21.43	0.834	0.390	22.52	0.796	0.341	22.43	0.796	0.258	27.51	0.965	0.245
Ours w/ object	21.56	0.767	0.269	23.98	0.909	0.327	21.08	0.830	0.396	22.14	0.787	0.347	21.88	0.785	0.269	26.80	0.962	0.246
Ours w/ state, object	22.33	0.785	0.253	24.08	0.909	0.320	21.51	0.833	0.386	22.55	0.796	0.338	22.38	0.796	0.259	27.52	0.965	0.245
Ours w/ state, object, mask	22.48	0.790	0.245	24.19	0.911	0.321	21.60	0.834	0.382	22.77	0.799	0.335	22.48	0.802	0.251	27.73	0.968	0.227
Ours w/ state, object, mask, cycle	22.61	0.795	0.241	24.19	0.911	0.321	21.71	0.835	0.382	22.68	0.802	0.336	22.54	0.804	0.249	27.88	0.968	0.227
Ours (full)	22.56	0.792	0.243	24.15	0.911	0.320	21.74	0.836	0.382	22.67	0.801	0.336	22.63	0.804	0.248	27.74	0.968	0.227

Table 2: Per-scene quantitative evaluation on the HOSNeRF dataset against baselines and ablations of our method. We color code each cell as **best** and **second best**.

	SEATTLE			PARKING			BIKE			JOGGING			LAB			CITRON		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
NSFF [20]	21.84	0.69	0.37	21.98	0.69	0.46	21.16	0.71	0.36	20.63	0.53	0.49	20.52	0.75	0.39	12.33	0.49	0.65
HyperNeRF [30]	16.43	0.43	0.40	16.04	0.38	0.62	17.64	0.42	0.43	18.52	0.39	0.52	16.75	0.51	0.23	16.81	0.41	0.56
NeuMan [14]	23.98	0.77	0.26	25.43	0.79	0.31	25.52	0.82	0.23	22.68	0.67	0.32	24.93	0.85	0.21	24.71	0.80	0.26
Ours	26.68	0.91	0.14	27.20	0.93	0.12	26.09	0.93	0.10	25.04	0.89	0.16	24.93	0.94	0.10	24.44	0.90	0.18

Table 3: Per-scene quantitative evaluation on the NeuMan dataset against baselines. We color code each cell as **best**.

ter foreground-background merging results, such as the feet regions of the Seattle scene and the Lab scene. This further demonstrates the effectiveness and flexibility of HOSNeRF to model various types of dynamic human-(objects)-scenes.

Per-Scene 360° Free-Viewpoint Renderings from Our HOSNeRF. We have also included a supplementary video to showcase the per-scene 360° free-viewpoint renderings from our HOSNeRF on all six scenes of our HOSNeRF dataset. The video highlights that our HOSNeRF is capable of generating high-fidelity details that accurately resemble all scene components, including dynamic human bodies, objects, and backgrounds. Notably, our HOSNeRF facilitates pausing the monocular video at any given point and rendering all scene details with high-fidelity from 360° viewpoints. To the best of our knowledge, our HOSNeRF represents the first work to achieve 360° free-viewpoint high-fidelity novel view synthesis for dynamic scenes featuring human-environment interactions from a single video.

It should be emphasized that to enable 360° free-viewpoint rendering, the input videos must contain 360° scene information (an example capture process is shown in the supplementary video); otherwise, artifacts will appear in unobserved areas that were not even seen in the input videos. However, our HOSNeRF is not restricted by such capturing requirements, as demonstrated on the NeuMan dataset [14], which features small camera motions.

4.4. Ablation Study

We conduct ablation studies on our collected dataset to evaluate the effectiveness of each proposed component in HOSNeRF. We progressively ablate each component from

state embeddings, object bones, foreground masks, optical flow, and cycle consistency supervision. As shown in Tab. 2, the performance of HOSNeRF progressively drops with the disabling of each component, with the full model nearly achieving the best performances, which demonstrates the effectiveness of our designs.

We visualize representative scenes with large (Fig. 6) and small (Fig. 7) objects to further evaluate our proposed designs. As shown in Fig. 6(b) and 7(b), the base models are unable to represent the interacted objects for the human and the scene. In contrast, the proposed state embedding is capable of representing the state transitions of small objects (Fig. 7(c)), but it fails to represent the deformations of large objects (Fig. 6(c)). Together with our proposed object bones, the deformation of large objects can be effectively estimated through the human-object skeleton hierarchy, enabling large objects modeling in Fig. 6(d). In addition, the full model can further improve the synthesis quality of fine details such as correcting the wrong right arm by comparing Fig. 7(c) and (d). These comparisons further demonstrate the importance and effectiveness of the proposed designs.

5. Conclusion

We introduced a novel framework of HOSNeRF, the first work to achieve 360° free-viewpoint high-fidelity novel view synthesis for dynamic scenes with human-environment interactions from a single video. To tackle the challenges of such a task, we first introduced the new object bones to the human skeleton hierarchy to effectively estimate the fast human-object motion in our dynamic human-object model. Then, we proposed the state-conditional



Figure 5: Qualitative comparisons of HOSNeRF and NeuMan [14] on the NeuMan dataset.

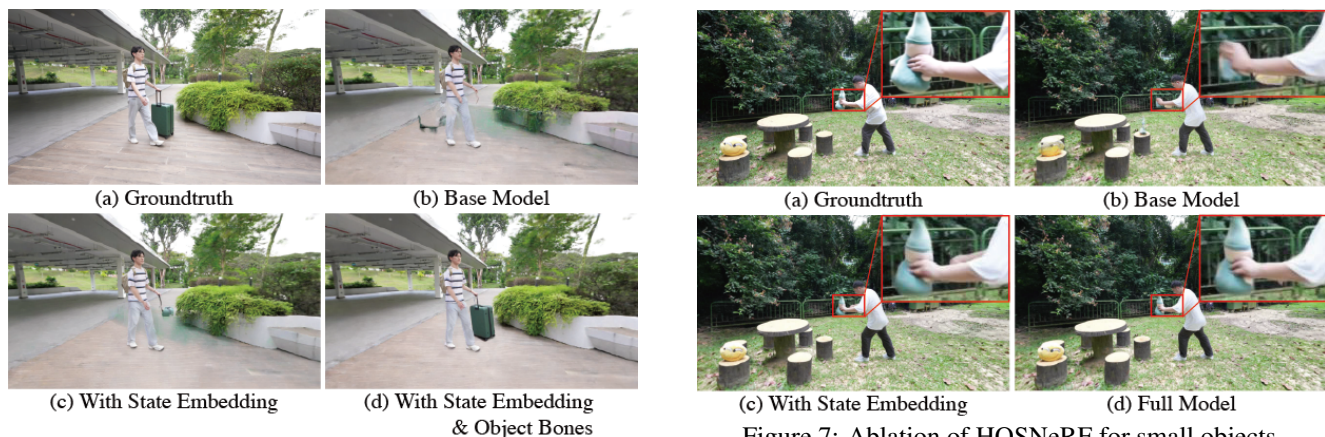


Figure 6: Ablation of HOSNeRF for large objects.

human-object and scene representations for handling significant object state changes. With our designed training strategies, HOSNeRF produced significant improvements over SOTA approaches, and enabled 360° free-viewpoint rendering of any frame from a single video and supported rendering views with novel human and object poses.

Limitations and Future Work. Our HOSNeRF focuses on dynamic human-objects modeling and lacks the capability to represent dynamic backgrounds. It remains challenging but is worthwhile researching models for reconstructing high-fidelity dynamic unbounded backgrounds.

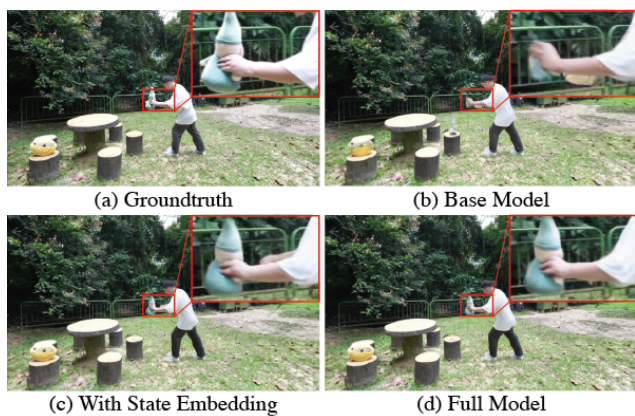


Figure 7: Ablation of HOSNeRF for small objects.

6. Acknowledgements

This research is supported by the National Research Foundation, Singapore under its NRFF Award NRF-NRFF13-2021-0008, Mike Zheng Shou’s Start-Up Grant from NUS, the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2022) Award 22-5406-A0001. This research is also supported by National Research Foundation, Singapore and A*STAR, under its RIE 2020 Industry Alignment Fund - Industry Collaboration Projects (IAF-ICP) grant call (grant no. I2001E0059). Jia-Wei Liu is also supported by NUS IDS-ISEP scholarship. Thanks to Yufei Shi for helpful discussions.

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 696–712. Springer, 2020.
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022.
- [3] Ang Cao and Justin Johnson. Hexplane: a fast representation for dynamic scenes. *arXiv preprint arXiv:2301.09632*, 2023.
- [4] Mingfei Chen, Jianfeng Zhang, Xiangyu Xu, Lijuan Liu, Yujun Cai, Jiashi Feng, and Shuicheng Yan. Geometry-guided progressive nerf for generalizable and efficient neural human rendering. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIII*, pages 222–239. Springer, 2022.
- [5] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (ToG)*, 34(4):1–13, 2015.
- [6] Jiemin Fang, Taoran Yi, Xingang Wang, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022.
- [7] Sara Fridovich-Keil, Giacomo Meanti, Frederik Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. *arXiv preprint arXiv:2301.10241*, 2023.
- [8] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5712–5721, 2021.
- [9] Xiangjun Gao, Jiaolong Yang, Jongyoo Kim, Sida Peng, Zicheng Liu, and Xin Tong. Mps-nerf: Generalizable 3d human rendering from multiview images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–12, 2022.
- [10] 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(4):1–16, 2021.
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [12] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3093–3102, 2020.
- [13] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. Volumedeform: Real-time volumetric non-rigid reconstruction. In *European Conference on Computer Vision*, pages 362–379. Springer, 2016.
- [14] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, pages 402–418. Springer, 2022.
- [15] Yuheng Jiang, Suyi Jiang, Guoxing Sun, Zhuo Su, Kaiwen Guo, Minye Wu, Jingyi Yu, and Lan Xu. Neuralhofusion: Neural volumetric rendering under human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6155–6165, 2022.
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [17] Youngjoong Kwon, Dahun Kim, Duygu Ceylan, and Henry Fuchs. Neural human performer: Learning generalizable radiance fields for human performance rendering. *Advances in Neural Information Processing Systems*, 34:24741–24752, 2021.
- [18] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhöfer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, pages 419–436. Springer, 2022.
- [19] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5521–5531, 2022.
- [20] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021.
- [21] Zhengqi Li, Qianqian Wang, Forrester Cole, Richard Tucker, and Noah Snavely. Dynibar: Neural dynamic image-based rendering. *arXiv preprint arXiv:2211.11082*, 2022.
- [22] Jia-Wei Liu, Yan-Pei Cao, Weijia Mao, Wenqiao Zhang, David Junhao Zhang, Jussi Keppo, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Devrf: Fast deformable voxel radiance fields for dynamic scenes. *arXiv preprint arXiv:2205.15723*, 2022.
- [23] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Transactions on Graphics (TOG)*, 40(6):1–16, 2021.
- [24] 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)*, 38(5):1–14, 2019.
- [25] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural vol-

- umes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019.
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015.
- [27] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [28] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 343–352, 2015.
- [29] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021.
- [30] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021.
- [31] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14314–14323, 2021.
- [32] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021.
- [33] Sergey Prokudin, Michael J Black, and Javier Romero. Smpix: Neural avatars from 3d human models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1810–1819, 2021.
- [34] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021.
- [35] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 84–93, 2020.
- [36] Ruizhi Shao, Hongwen Zhang, He Zhang, Mingjia Chen, Yan-Pei Cao, Tao Yu, and Yebin Liu. Doublefield: Bridging the neural surface and radiance fields for high-fidelity human reconstruction and rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15872–15882, 2022.
- [37] Qing Shuai, Chen Geng, Qi Fang, Sida Peng, Wenhao Shen, Xiaowei Zhou, and Hujun Bao. Novel view synthesis of human interactions from sparse multi-view videos. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10, 2022.
- [38] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [39] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019.
- [40] Liangchen Song, Anpei Chen, Zhong Li, Zhang Chen, Lele Chen, Junsong Yuan, Yi Xu, and Andreas Geiger. Nerfplayer: A streamable dynamic scene representation with decomposed neural radiance fields. *arXiv preprint arXiv:2210.15947*, 2022.
- [41] Timo Stich, Christian Linz, Georgia Albuquerque, and Marcus Magnor. View and time interpolation in image space. In *Computer Graphics Forum*, volume 27, pages 1781–1787. Wiley Online Library, 2008.
- [42] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. *Advances in Neural Information Processing Systems*, 34:12278–12291, 2021.
- [43] Zhuo Su, Lan Xu, Dawei Zhong, Zhong Li, Fan Deng, Shuxue Quan, and Lu Fang. Robustfusion: Robust volumetric performance reconstruction under human-object interactions from monocular rgbd stream. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [44] Guoxing Sun, Xin Chen, Yizhang Chen, Anqi Pang, Pei Lin, Yuheng Jiang, Lan Xu, Jingyi Yu, and Jingya Wang. Neural free-viewpoint performance rendering under complex human-object interactions. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4651–4660, 2021.
- [45] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Michael J Black, and Tao Mei. Monocular, one-stage, regression of multiple 3d people. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11179–11188, 2021.
- [46] Xin Suo, Yuheng Jiang, Pei Lin, Yingliang Zhang, Minye Wu, Kaiwen Guo, and Lan Xu. Neuralhumanfvv: Real-time neural volumetric human performance rendering using rgb cameras. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6226–6237, 2021.
- [47] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12959–12970, 2021.
- [48] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Fourier plenotrees for dynamic radiance field rendering in real-time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13524–13534, 2022.

- [49] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [50] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16210–16220, 2022.
- [51] Minye Wu, Yuehao Wang, Qiang Hu, and Jingyi Yu. Multi-view neural human rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1682–1691, 2020.
- [52] Tianhao Wu, Fangcheng Zhong, Andrea Tagliasacchi, Forrester Cole, and Cengiz Oztireli. D²nerf: Self-supervised decoupling of dynamic and static objects from a monocular video. *Advances in Neural Information Processing Systems*, 35:32653–32666, 2022.
- [53] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9421–9431, 2021.
- [54] Hongyi Xu, Thiemo Alldieck, and Cristian Sminchisescu. H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. *Advances in Neural Information Processing Systems*, 34:14955–14966, 2021.
- [55] Jiakai Zhang, Xinhang Liu, Xinyi Ye, Fuqiang Zhao, Yanshun Zhang, Minye Wu, Yingliang Zhang, Lan Xu, and Jingyi Yu. Editable free-viewpoint video using a layered neural representation. *ACM Transactions on Graphics (TOG)*, 40(4):1–18, 2021.
- [56] Juze Zhang, Haimin Luo, Hongdi Yang, Xinru Xu, Qianyang Wu, Ye Shi, Jingyi Yu, Lan Xu, and Jingya Wang. Neuraldome: A neural modeling pipeline on multi-view human-object interactions. *arXiv preprint arXiv:2212.07626*, 2022.
- [57] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [58] C Lawrence Zitnick, Sing Bing Kang, Matthew Uyttendaele, Simon Winder, and Richard Szeliski. High-quality video view interpolation using a layered representation. *ACM transactions on graphics (TOG)*, 23(3):600–608, 2004.
- [59] Michael Zollhöfer, Matthias Nießner, Shahram Izadi, Christoph Rehmann, Christopher Zach, Matthew Fisher, Chenglei Wu, Andrew Fitzgibbon, Charles Loop, Christian Theobalt, et al. Real-time non-rigid reconstruction using an rgb-d camera. *ACM Transactions on Graphics (ToG)*, 33(4):1–12, 2014.