

Tensor4D : Efficient Neural 4D Decomposition for High-fidelity Dynamic Reconstruction and Rendering

Ruizhi Shao¹, Zerong Zheng^{1,2}, Hanzhang Tu¹, Boning Liu¹, Hongwen Zhang¹, Yebin Liu¹
¹Department of Automation, Tsinghua University ²NNKosmos Technology




Figure 1. Given 4 sparse static RGB camera views of a dynamic scene (a), our proposed Tensor4D decomposition enables multiview reconstruction to achieve fine-grained geometry reconstruction even on human fingers (b) and temporal-consistent novel view synthesis on a 3D holographic display (c,d,e). The 4 cameras are settled on four corners of the display. The proposed method demonstrates low-cost, portable and highly immersive telepresence experience.

Abstract

We present *Tensor4D*, an efficient yet effective approach to dynamic scene modeling. The key of our solution is an efficient 4D tensor decomposition method so that the dynamic scene can be directly represented as a 4D spatio-temporal tensor. To tackle the accompanying memory issue, we decompose the 4D tensor hierarchically by projecting it first into three time-aware volumes and then nine compact feature planes. In this way, spatial information over time can be simultaneously captured in a compact and memory-efficient manner. When applying *Tensor4D* for dynamic scene reconstruction and rendering, we further factorize the 4D fields to different scales in the sense that structural motions and dynamic detailed changes can be learned from coarse to fine. The effectiveness of our method is validated on both synthetic and real-world scenes. Extensive experiments show that our method is able to achieve high-quality dynamic reconstruction and rendering from sparse-view camera rigs or even a monocular camera. The code and dataset will be released at <https://github.com/DSaurus/Tensor4D>.

1. Introduction

High quality reconstruction and Photo-realistic rendering of a dynamic scene from a set of input images is nec-

essary for many applications such as AR/VR, 3D content production and entertainment. Traditional methods use classical mesh-based representation to reconstruct the dynamic scenes, which, unfortunately, are prone to produce reconstruction errors and rendering artifacts when the scene contains thin structures, specular surfaces and topological changes [9, 19, 23, 26, 49].

Recent advances in neural rendering approaches, which learn scene representations in the form of neural radiance fields (NeRF), have shown impressive novel view synthesis of general static scenes given only multi-view images [32]. They are immediately extended to dynamic scenes: some methods (e.g., NeRF-T) consider time as an additional input dimension to NeRF representation [56, 62], while other methods (e.g., D-NeRF) disentangle a dynamic scene into a canonical radiance field and a dynamic motion field [11, 28, 37, 40, 55]. Either way, learning a 4D function is one of the main cornerstones. Unfortunately, directly using MLP to fit such a function often suffers from high time and computation cost, i.e., dozens of hours on high-end GPUs.

In fact, the aforementioned limitation also exists in conventional NeRF-based methods for static scenes, and researchers have proposed to use discrete data structures like voxel grids [65] or triplanes [7] to accelerate NeRF training and rendering. However, these techniques are difficult to be extended to dynamic domains as introducing an additional time dimension will exponentially increase memory footprint, hindering them from modeling high-quality ap-

pearance details.

In this work, we pursue a dynamic scene representation that also utilizes explicit feature grids to accelerate network training while avoiding huge memory consumption when introducing an additional time dimension. To this end, we bypass the construction of a high resolution 4D tensor; instead, we propose to model a 4D field using hierarchical tri-projection decomposition. Our decomposition method extends the tri-projection in EG3D [7]. It firstly project a full 4D field into three time-aware volumes, each of which is then further decomposed into three feature planes. In this way, we model the 4D field using only nine 2D feature planes, and we empirically find that although being highly compact, such a representation is powerful enough to represent dynamic scenes containing complex motions. Moreover, the usage of explicit data structure also allows us to design a coarse-to-fine strategy to further improve the performance of our method.

By utilizing and factorizing an explicit 4D tensor, our method enables both efficient reconstruction and compact representation of dynamic scenes. Besides, the decomposition scheme also introduces implicit constraints on the representation since only low-rank tensors can be approximated by a small number of lower-dimensional components. Such a constraint can serve as an inherent regularization when the input observation is limited, *e.g.*, under sparse and fixed cameras setting or even monocular inputs. In this paper, we first apply our method for sparse-view dynamic reconstruction by adopting our Tensor4D decomposition to time-conditioned radiance fields in “NeRF-T”. In addition, our decomposition method can also be used for single-view dynamic reconstruction. This is achieved through decomposing both the 4D dynamic motion field and the canonical radiance field in “D-NeRF”. With proper regularization, our system enables efficient and high-quality reconstruction of dynamic objects under both camera settings.

2. Related Work

Multiview Reconstruction and Rendering. There are various ways to capture and reconstruct the 3D dynamic scenes, involving methods based on silhouette [24, 54], stereo [30, 46], flow [17, 44, 70], segmentation [41, 42], and photometric [3, 20]. With RGBD cameras, real-time solutions like DynamicFusion [34] estimates the non-rigid deformations of a dynamic scene and integrates depth frames to reconstruct the geometry model in the canonical space. This method is later extended for telepresence and holographic communication [10, 25, 66]. However, these systems heavily rely on depth sensors to obtain accurate geometry. In contrast, our method can reconstruct and render dynamic scenes using sparse-view RGB cameras.

In the past few years, neural implicit representations underwent rapid development and have been applied for multi-

view reconstruction and rendering of static scenes. Some methods represent the geometry as the zero level-set of a neural network, and use differentiable surface rendering to optimize the network weights [35, 47, 64]. Given dense observation of an object, these methods are able to accurately recover its surface. NeRF [32], on the other hand, uses volume rendering to optimize the scene representation. Its simplicity and impressive results inspires a lot of following works, including in-the-wild reconstruction [31, 52], lighting and material estimation [4, 48, 67], generation [1, 6, 18, 22, 43, 53], human rendering [45, 51, 57, 68, 69] and so on. More recently, several methods unify surface and volumetric rendering, enabling accurate geometry reconstruction and high-quality novel view synthesis [36, 59, 63]. Compared to these static scene representations, we aim to enable free view synthesis of dynamic scenes from extremely sparse cameras.

NeRF for Dynamic Scenes. Modeling scenes in 4D domain with time dimension included is a direct solution to extend NeRF for dynamic domains. Typical approaches includes VideoNeRF [62], NeRFlow [11], DyNeRF [27], and DCT-NeRF [56]. Specifically, VideoNeRF [62] directly learns a spatiotemporal irradiance field from a single video and uses depth estimation to address the shape-motion ambiguities in monocular inputs, while NeRFlow [11] and DCT-NeRF [56] use point trajectory to regularize the network optimization. To deal with the limitation of topology-change modeling in deformation field, Park *et al.* [38] represent HyperNeRF which can lift NeRF to higher dimensions.

Dynamic scenes can also be rendered by deforming the radiance field in the canonical space. For example, Nerfies [37] optimizes an additional continuous deformation field by warping each observed point into a canonical 5D NeRF. D-Nerf [40] and NR-Nerf [55] follow a similar framework, but take only monocular videos as training data. In addition, DeVRF [28] uses voxel-based representation instead of MLPs to model the 3D canonical space and the 4D deformation field. Using parametric body templates as the semantic prior, methods like Neural Body [39] and Human-NeRF [60] enable photo-realistic novel view synthesis of complex human performance. The discrepancy between the practical capture process and the existing experimental protocols for monocular videos has been shown in [16].

NeRF Acceleration. Numerous works emerge with the purpose of speeding up static NeRF using explicit data structures including feature maps, voxels and tensors. For instance, DVGO [50] achieves fast convergence through an explicit representation of a density voxel grid and a feature voxel grid. With the sparse voxel octree structure, NSVF [29] accelerates the novel view reconstruction by discarding the empty voxels in a coarse-to-fine manner. Similarly, Plenoxels [14] and PlenOctree [65] model a scene through a hierarchical 3D grid with spherical harmonics, which can realize an optimization with two orders of magnitude faster

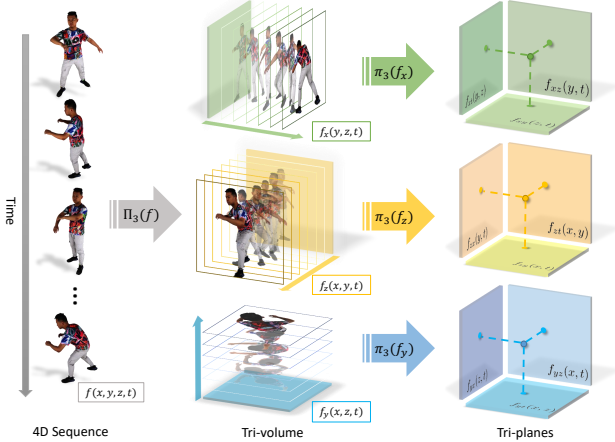


Figure 2. **Illustration of our hierarchical tri-projection decomposition method.** For a neural 4D field $f(x, y, z, t)$, we first decompose the 3D space part from 4D spatio-temporal tensor into three time-aware volumes, which are then further projected onto nine 2D planes.

than NeRF. In DIVER [61], ray marching only finds a fixed number of hits on the voxel grid to accelerate volumetric rendering. Moreover, hashing encoding [33] and tensor decomposition [8] are also used as compact representations for NeRF acceleration.

For dynamic scene modeling, DeVRF [28] enables fast non-rigid neural rendering with both 3D volumetric and 4D voxel field. In addition, V4D [15] introduce an effective conditional positional encoding for 4D data to realize fast novel view synthesis. TiNeuVox [12] represents scenes with optimizable explicit data structures and accelerates radiance fields modeling, while Wang *et al.* extended PlenOc-trees [65] into free-view video rendering [58]. However, the low-res volumetric design of these works hinders the capacity of rendering high-quality images. In addition, there are also several concurrent works [5, 13, 21] adopt 6-plane decomposition for dynamic scenes. Compared with these methods, our hierarchical decomposition is more efficient to capture time variations and represent dynamic scenes.

3. Method

Building on prior work for spatio-temporal representations and tri-projection decomposition (Sec. 3.1), we propose a hierarchical tri-projection decomposition method (Sec. 3.2) and a coarse-to-fine strategy (Sec. 3.3), which allow us to learn a 4D field at a modest cost of training time and GPU memory.

3.1. Preliminary

Spatio-temporal 4D NeRF representation. To represent dynamic objects using neural radiance fields, a naive way is to condition the original neural radiance fields on the timestamp [56, 62], which we term as NeRF-T. Mathematically,

NeRF-T can be formulated as:

$$\begin{aligned} f(x, y, z, t) &= (\mathbf{f}, \sigma), \\ g(\mathbf{f}, \theta, \phi) &= \mathbf{c}, \end{aligned} \quad (1)$$

where f is a dynamic implicit field that produces a high-dimensional feature \mathbf{f} and a density value σ for a point at position (x, y, z) and time instance t , and g is a function that takes into account the viewing direction (θ, ϕ) to predict the final RGB color.

To achieve better disentanglement of shape and motion, some methods like D-NeRF [40] propose deformable neural radiance field which adopts a canonical 3D representation with the 4D flow fields:

$$\begin{aligned} f(x, y, z, t) &= (\hat{x}, \hat{y}, \hat{z}), \\ g(\hat{x}, \hat{y}, \hat{z}, \theta, \phi) &= (\mathbf{c}, \sigma), \end{aligned} \quad (2)$$

where g is the radiance field in canonical configuration and f is a scene flow field representing the mapping between the scene at time instant t and the canonical space.

From the above formulation one can easily observe that both NeRF-T and D-NeRF rely on modeling a 4D field $f(x, y, z, t)$, *i.e.*, the dynamic implicit field in NeRF-T and the flow field in D-NeRF. Existing methods mainly adopt MLPs to fit these 4D fields. Such an implicit neural representation does not have an explicit structure and requires extensive computation time for both training and rendering. **Tri-projection Decomposition.** The tri-projection decomposition is widely used in recent work including EG3D [7] and TensorRF [8] in order to accelerate the training and rendering process in MLP-based NeRF frameworks. Such decomposition factorizes a n -dimensional tensor V_h into three lower-dimensional $((n-1)$ -D) tensors $V_l^i \{i=1, 2, 3\}$ by projecting V_h along the x , y and z -axis respectively. For example, triplane-based decomposition proposed by EG3D projects the 3D tensor into three 2D feature planes. Compared to voxel-based representations, triplane representation effectively reduces the memory footprint and improves the performance of 3D generation and reconstruction.

3.2. Hierarchical Tri-projection Decomposition

Our goal is to design a dynamic scene representation with explicit feature grids to accelerate network training and volumetric rendering. However, directly constructing a 4D tensor costs a huge amount of memory and is unacceptable for the purpose of high-resolution rendering. Therefore, we propose a hierarchical triprojection decomposition to factorize the 4D tensor into several compact features, which reduces memory consumption by a large margin while preserving the capability of fitting 4D fields.

Specifically, for a neural 4D field $f(x, y, z, t)$, we first decompose the 3D space part from 4D spatio-temporal tensor into three time-aware volumetric tensors via tri-

projection decomposition:

$$\Pi_3(f(x, y, z, t)) = \{f_z(x, y, t), f_y(x, z, t), f_x(y, z, t)\}, \quad (3)$$

where Π_3 denotes the projection operator. To further lower space complexity and enable high-resolution representation, we decompose each feature volume into three feature planes as:

$$\begin{aligned} \pi_3(f_z) &= \{f_{zt}(x, y), f_{zy}(x, t), f_{zx}(y, t)\} \\ \pi_3(f_y) &= \{f_{yt}(x, z), f_{yz}(x, t), f_{yx}(z, t)\} \\ \pi_3(f_x) &= \{f_{xt}(y, z), f_{xz}(y, t), f_{xy}(z, t)\} \end{aligned} \quad (4)$$

where π_3 denotes the volume-to-plane tri-projection. In this way, we compactly represent a 4D field using 9 planes. Given any spatio-temporal coordinate (x, y, z, t) , we can efficiently query its value in the 4D field by projecting it onto the planes and retrieving the corresponding value via bilinear interpolation. Fig 2 is an illustration of our decomposition method. Our hierarchical triprojection decomposition reduces the space complexity from $\mathcal{O}(n^4)$ to $\mathcal{O}(n^2)$ with n being the spatial resolution of the grid, significantly lowering memory footprint without sacrificing representation power.

Differences from 6-plane decomposition. Compared with 6-plane decomposition [5, 13, 21] which pairs the time dimension with only one spatial dimension (xt, yt, zt) , our method first decomposes the 4D tensor in the spatial domain to obtain three time-aware volumes and then 9 decomposed planes. Note that the three volumes are decomposed independently, making the 9 planes different from each other. Specifically, when using a linear layer $g = \{\cdot\}$ to decompose the volume $f_z(x, y, t) = \{\cdot, f_{zy}(x, t), \cdot\}$ and the volume $f_y(x, z, t) = \{\cdot, f_{yz}(x, t), \cdot\}$, the planes $f_{zy}(x, t)$ and $f_{yz}(x, t)$ behave differently since they are optimized independently to capture time variations in the volumes $f_z(x, y, t)$ and $f_y(x, z, t)$, respectively. In this way, the 9 planes can leverage time-aware information in all possible combinations of the time and spatial dimensions $(xt, yt, zt, xyt, yzt, xzt, xyz)$. This property enables our representation to capture various dynamic information hierarchically at different levels of spatial dimensions. Therefore, our method is more efficient in representing dynamic scenes, where variations in time dimension are often intense, complex, and long-range.

Our decomposition supports various types of 4D fields including the time-conditioned radiance field in NeRF-T and the 4D flow field in D-NeRF. In this paper, we present a dynamic radiance field decomposition method (Sec. 4.1) for multiview dynamic reconstruction as well as a 4D flow decomposition method (Sec. 4.2) for single-view setting.

3.3. Coarse-to-fine Strategy

To further improve the efficiency of our 4D decomposition, we propose an optional coarse-to-fine strategy to fac-

torize the 4D fields into different scales in different training phases. In coarse level, we adopt low-resolution feature planes (128×128) to decompose the 4D fields, which improves the robustness of the training process and achieves fast convergence. After coarse level training, we additionally use high-resolution feature planes (512×512) for 4D decomposition to represent dynamic details and achieve high-quality rendering. Specifically, we factorize the 4D fields into different scales:

$$f(x, y, z, t) = \{\pi_3(\Pi_3(f^{LR})), \pi_3(\Pi_3(f^{HR}))\}, \quad (5)$$

where f^{LR} and f^{HR} are the coarse-level and fine-level components of $f(x, y, z, t)$, respectively. In the coarse level, the decomposed feature planes is low-resolution to represent coarse 3D structures and 4D dynamic changes. In the fine level, we adopt the high-resolution feature planes in each element to decompose the 4D fields, which focuses more on recovering dynamic details.

4. Tensor4D for Dynamic Reconstruction

As shown in Fig. 3, we apply our 4D tensor decomposition into the task of dynamic reconstruction with two inputs:

1) Dynamic reconstruction under sparse and fixed cameras. For this setting, we factorize NeRF-T instead of D-NeRF using our proposed 4D decomposition, which we found more efficient and flexible to represent topologically-varying objects (Sec. 4.1).

2) Dynamic reconstruction using monocular camera. For this setting, we separately decompose the 4D flow fields and the 3D canonical representation since it ensures appearance consistency across different frames and achieves more robust performance (Sec. 4.2).

4.1. Multi-view Reconstruction

In this section, we present our dynamic reconstruction system under sparse multi-view setting. The goal of our system is 1) efficient and high-quality dynamic reconstruction with low memory and time cost, and 2) robust reconstruction under sparse and fixed camera setting.

To this end, we adopt our 4D decomposition (Eq. 5) to factorize the NeRF-T representation with the coarse-to-fine strategy. Specifically, we can obtain nine low-resolution feature planes T^{LR} and nine high-resolution feature planes T^{HR} to represent the 4D NeRF-T fields after our 4D decomposition. Then for volume rendering, when sampling a point p at (x, y, z, t) in the ray r with direction of (θ, ϕ) , we first query the feature F_p of p in both the LR and the HR feature planes. Take the LR features for example:

$$\begin{aligned} F_p^{LR}(x, y, z, t) &= \pi_3(\Pi_3(f^{LR})) \\ &= \oplus(\pi_3(f_z^{LR}), \pi_3(f_y^{LR}), \pi_3(f_x^{LR})) \\ &= \oplus(\oplus(T_{zt}^{LR}(x, y), T_{zy}^{LR}(x, t), T_{zx}^{LR}(y, t)), \dots), \end{aligned} \quad (6)$$




Figure 3. **The framework of Tensor4D for multi-view and monocular reconstruction.** **a).** Tensor4D for multi-view reconstruction. The 4D NeRF-T fields are separately factorized by the LR and HR feature planes. **b).** Tensor4D for monocular reconstruction. The 4D flow fields are factorized by the LR feature plane for better disentanglement of shape and motion. The 3D canonical representation is factorized by three LR and HR feature planes.

where we concatenate all features queried in the nine LR planes as the final LR feature of the sampling point p . Then we concatenate the LR and HR features with the positional encoding of (x, y, z, t) into the geometry MLP E_g to obtain the density σ and high-dimension feature \mathbf{f} :

$$E_g(F_p^{LR}, F_p^{HR}, \gamma(x, y, z, t)) = (\mathbf{f}, \sigma), \quad (7)$$

where $\gamma()$ is the positional encoding function. Next, we concatenate the high-dimension feature \mathbf{f} with the positional encoding of (θ, ϕ) and feed it into the color MLP:

$$E_c(\mathbf{f}, \gamma(\theta, \phi)) = c \quad (8)$$

In this way, we can render the images through volume rendering and adopt color loss to train our decomposed feature planes T and neural network E_c and E_g :

$$L_c = \|C_{i,j} - C_{i,j}^*\| \quad (9)$$

With our coarse-to-fine design, our method can recover high-fidelity dynamic details effectively and efficiently.

To achieve robust dynamic reconstruction under sparse multi-view setting, we further adopt regularization for all decomposed feature planes T :

$$L_r = \sum_T \sum_{i,j} \sqrt{(T_{i+1,j} - T_{i,j})^2 + (T_{i,j+1} - T_{i,j})^2}, \quad (10)$$

where L_r is the TV loss for each feature plane T to keep their sparsity. To regularize the geometry, we also introduce surface constraint into volume rendering. Specifically, we adopt the SDF as the base geometry representation and follow NeuS [59] to render the SDF field. Then we add surface constraint loss to enforce a smooth surface:

$$L_e = \| \|\nabla s(x, y, z, t)\|_2 - 1 \|_2 \quad (11)$$

The final training loss L_m is the regularization loss L_r , L_e and the color loss L_c :

$$L_m = \lambda_r L_r + \lambda_e L_e + \lambda_c L_c \quad (12)$$

4.2. Monocular Reconstruction

Different from the sparse view setting, we adopt 4D decomposition for D-NeRF in monocular capture cases. This is because monocular setting is much more ill-posed than sparse-view inputs, and the explicit disentanglement of appearance and motions can guarantee the consistency across different frames. Since D-NeRF represents the dynamic objects with the 4D flow fields and a 3D canonical representation, we separately factorize these two fields. First, for the 4D flow fields, we only adopt coarse level decomposition and factorize it into low-resolution feature planes:

$$f(x, y, z, t) = \pi_3(\Pi_3(f^{LR})). \quad (13)$$

Our coarse decomposition focuses more on the coarse and rigid motion, which improves the robustness of flow estimation and can achieve better disentanglement of shape and motion. Then for the 3D canonical representation, we adopt both coarse and fine level decomposition:

$$h(\hat{x}, \hat{y}, \hat{z}) = \pi_3(h^{HR}) + \pi_3(h^{LR}). \quad (14)$$

Therefore, we can obtain 9 flow feature planes T_f for 4D flow fields and 6 canonical feature planes T_h for 3D canonical representation. The volume rendering in monocular cases is similar to multi-view cases. For a sampling point p , we first obtain the point flow feature F_f using Eq. 6 with nine flow planes T_f . Then we adopt the flow MLP E_f to predict the movement of the point:

$$E_f(F_p, \gamma(x, y, z, t)) = (\hat{x}, \hat{y}, \hat{z}). \quad (15)$$

Then we obtain the point canonical feature F_c by querying the canonical feature planes T_h . Take the LR feature F_c^{LR} for example:

$$F_c^{LR} = \pi_3(h^{LR}) = \oplus (T_{h_x}^{LR}(x, y), T_{h_y}^{LR}(x, z), T_{h_z}^{LR}(x, y)). \quad (16)$$

Next, we feed canonical feature F_c^{LR} and the positional encoding of $(\hat{x}, \hat{y}, \hat{z})$ into the geometry MLP E_g to obtain




Figure 4. **Example results of our method.** Space and time novel view rendering results from sparse-view fixed cameras. The top three results are from four front view cameras and the bottom is from 12 circular cameras.

high-dimension feature \mathbf{f} and density σ :

$$E_g(F_c^{LR}, \gamma(\hat{x}, \hat{y}, \hat{z})) = (\mathbf{f}, \sigma). \quad (17)$$

Finally, we adopt Eq. 8 to predict color c for volume rendering and Eq. 9 for training color loss L_c . We also add feature regularization loss L_r in Eq. 10 and surface constraint loss L_e in Eq. 11. The total training loss L_s is:

$$L_s = \lambda_c L_c + \lambda_r L_r + \lambda_e L_e. \quad (18)$$

5. Experiment

Dataset. To evaluate the performance of our methods for multiview inputs, we build a sparse-view capture system with 6 forward-facing RGB cameras mounted on the borders of a 32” Looking Glass 3D holographic display [2]. All cameras are synchronized and calibrated. Using this system, we capture multiple sequences of various challenging human motions, including dancing, thumbing up, waving hands, wearing hats and manipulating bags. We use 4 of them for reconstruction and rendering in all our experiments, while leaving the other two for quantitative evaluation. We also use three 360° multiview full body sequences

captured with 12 evenly spaced cameras on a camera ring for qualitative evaluation. For monocular evaluation, we use the synthetic dataset provided by D-NeRF [40] and select 3 scenes (“lego”, “standup” and “jumpingjacks”) from this dataset, with the numbers of training frames ranging from 50 to 200. More details about data collection and preprocessing can be found in the *Supp.Mat.*.

Baselines. We mainly compare our method against the following state-of-the-art baselines that are most related to our work: D-NeRF [40], NeRF-T, TiNeuVox [12] and NeuS-T. Here, NeRF-T is our extension of vanilla NeRF [32] by introducing an additional time input, and NeuS-T is extended from NeuS [59] similarly. Among these baselines, D-NeRF and TiNeuVox represent the dynamic scenes through deforming a canonical one, while NeRF-T and NeuS-T directly learn a time-conditioned 4D radiance fields. TiNeuVox uses explicit voxel grids to accelerate network training, while others purely use MLPs to model the scene.

5.1. Results and Comparison

Qualitative Results. We train our model for each individual sequences, and present some example results for novel view synthesis in Fig. 4 and the *Supp.Mat.*. The results




Figure 5. **Comparison on monocular synthetic dataset** against D-NeRF [40] and TiNeuVox [12].

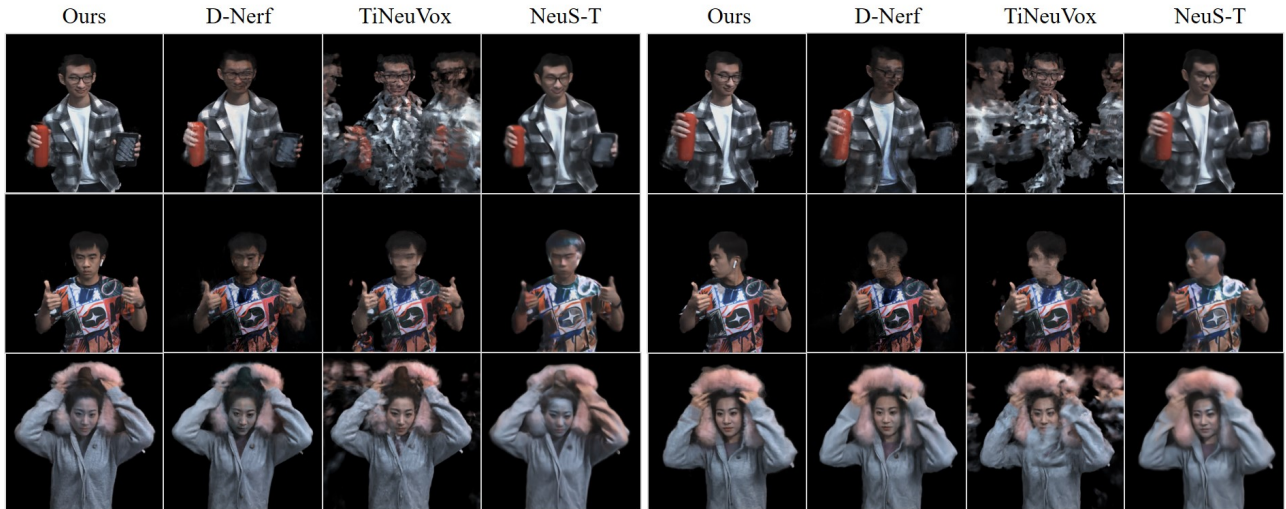


Figure 6. **Comparison on sparse-view real-world dataset** against D-NeRF [40], TiNeuVox [12] and NeuS-T [59]. Four of camera views from our datasets are used for reconstruction and rendering.

cover various body motions, clothing styles and accessories. As shown in Fig. 4, our method can render high-quality images for dynamic scenes and faithfully recover appearance details like the thin finger motions, semi-transparent silk, hand-object interaction, face expressions and cloth wrinkles. See our *Supp.Video.* for better visualization.

Comparisons on monocular dynamic dataset. We first compare our method with the baselines on monocular synthetic dataset. Qualitative results are presented in Fig. 5. Compared to other methods, ours recovers more appearance details and generate less artifacts. The numeric results in Tab. 1 also prove that our method outperforms state-of-the-art methods in terms of rendering quality and accuracy.

Comparisons on sparse view dataset. We then evaluate the performance of different methods for novel view synthesis given four camera views from our collected six camera-view real-world datasets. The remained two views are used

for quantitative evaluation. Results reported in Fig. 6 and Tab. 2 show again that our method performs better in accurate and high-quality appearance detail synthesis.

Comparisons of training efficiency. We compare the training time and model size for memory in Tabs. 3 and 4. Our method requires significantly less training time compared to the original D-NeRF, NeRF-T, and NeuS-T in monocular and multi-view scenarios. Compared to Ours-NeRF-T and Ours-D-NeRF, the longer training time observed in our method is due to the additional computation for geometry smoothness terms, which reduces artifacts and enhances rendering quality. Compared to TiNeuVox, our method has a reduction in memory consumption and a clear improvement in rendering quality since Tensor4D enables the efficient decomposition of 4D fields at higher spatial resolutions (512^3 vs. 256^3).

Table 1. Quantitative comparisons on monocular synthetic datasets.

Method	Lego				Standup				Jumpingjacks			
	MSE↓	PSNR↑	SSIM↑	LPIPS↓	MSE↓	PSNR↑	SSIM↑	LPIPS↓	MSE↓	PSNR↑	SSIM↑	LPIPS↓
D-NeRF [40]	7.83e-3	21.26	0.869	1.72e-1	4.63e-4	34.38	0.989	2.18e-2	5.46e-4	33.37	0.987	4.84e-2
NeRF-T	3.89e-3	24.32	0.904	1.55e-1	6.82e-4	31.44	0.968	2.36e-2	5.99e-4	32.27	0.979	5.37e-2
TiNeuVox [12]	3.15e-3	25.14	0.924	8.37e-2	2.90e-4	36.18	0.986	2.02e-2	3.89e-4	34.76	0.983	3.33e-2
Ours-NeRF-T	3.10e-3	25.13	0.922	1.24e-1	5.92e-4	32.38	0.977	2.29e-2	5.62e-4	32.73	0.980	5.16e-2
Ours-D-NeRF	4.61e-3	23.37	0.890	1.19e-1	2.82e-4	35.93	0.981	2.07e-2	4.22e-4	34.10	0.982	3.41e-2
Ours	2.26e-3	26.71	0.953	3.49e-2	2.50e-4	36.32	0.983	1.74e-2	3.91e-4	34.43	0.982	3.18e-2

Table 2. Quantitative comparisons on the six camera-view real-world datasets.

Method	Sequence1-thz				Earphone				Sequence3-yxd			
	MSE↓	PSNR↑	SSIM↑	LPIPS↓	MSE↓	PSNR↑	SSIM↑	LPIPS↓	MSE↓	PSNR↑	SSIM↑	LPIPS↓
D-NeRF [40]	3.13e-3	25.15	0.910	0.185	8.16e-3	20.92	0.854	0.256	4.84e-3	23.22	0.937	0.147
TiNeuVox [12]	5.72e-3	22.79	0.832	0.209	1.49e-2	18.55	0.707	0.319	8.18e-3	21.19	0.816	0.233
Neus-T [59]	4.95e-3	23.07	0.887	0.130	3.75e-3	24.22	0.877	0.184	2.20e-3	26.59	0.945	0.099
Ours-NeRF-T	5.85e-3	22.46	0.842	0.191	1.27e-2	19.13	0.838	0.218	5.02e-3	23.06	0.914	0.141
Ours-D-NeRF	4.74e-3	23.27	0.864	0.176	8.17e-3	21.07	0.883	0.198	4.06e-3	23.98	0.926	0.130
Ours	1.53e-3	28.27	0.942	0.084	3.20e-3	25.00	0.903	0.153	1.31e-3	28.83	0.962	0.072

Table 3. Training time and memory report in monocular cases.

Method	Lego			
	PSNR	Time	Iterations	#Params
TiNeuVox	25.14	34min	50k	102M
D-NeRF	21.26	45h	800k	4.8M
Ours-D-NeRF	23.37	95min	50k	20M
NeRF-T	24.32	38h	800k	4.4M
Ours-NeRF-T	25.13	73min	50k	43M
Ours(6-planes)	26.34	135min	50k	17M
Ours	26.71	144min	50k	20M

Table 4. Training time and memory report in multi-view cases.

Method	Sequence1-thz			
	PSNR	Time	Iterations	#Params
TiNeuVox	22.79	30min	50k	102M
D-NeRF	25.15	37h	800k	4.8M
Neus-T	23.07	45h	800k	5.0M
Ours-NeRF-T	22.46	68min	50k	43M
Ours(6-planes)	27.86	105min	50k	32M
Ours	28.27	117min	50K	43M

Table 5. Ablation study of hierarchical decomposition and smoothness terms.

Method	Lego		Standup	
	PSNR	SSIM	PSNR	SSIM
Ours(w/o regular)	26.49	0.946	35.91	0.977
Ours(6-planes)	26.44	0.944	35.79	0.978
Ours	26.71	0.953	36.32	0.983
Method	Sequence-thz		Sequence-earphone	
	PSNR	SSIM	PSNR	SSIM
Ours(w/o regular)	27.92	0.932	24.75	0.885
Ours(6-planes)	27.74	0.934	24.39	0.889
Ours	28.27	0.942	25.00	0.903

5.2. Ablation study

Regularization. We quantitatively ablate the regularization terms in our method. We implement two strong baselines ‘‘Ours-NeRF-T’’ and ‘‘Ours-D-NeRF’’, in which we directly apply our Tensor4D decomposition for NeRF-T and D-NeRF without the regularization terms. The results are reported in Tab. 1 and Tab. 2. Benefiting from our 4D decomposition, they achieve better performance than the original NeRF-T and D-NeRF. However, their rendering quality is still worse than our full method with regularization. In

addition, We ablate for TV regular smoothness terms (Ours vs. Ours(w/o regular)) and the results are reported in Tab. 5, which further validates the effectiveness of our smoothness terms for rendering quality enhancement.

Hierarchical Decomposition We quantitatively ablate our hierarchical decomposition with 6-plane decomposition. As shown in Tab 5, our method achieves superior performance in both multi-view and monocular cases, which validates the effectiveness of our hierarchical decomposition.

6. Discussion

Limitations. Since our method needs to decompose the 4D fields to several 2D feature planes, a pre-set bounding box of the scene is necessary. Therefore, it is difficult for our method to reconstruct backgrounds or objects which are out of the bounding box. Another limitation is our strong regularization terms. Though these terms benefit the robustness of our reconstruction under sparse views, they also limit our ability to handle challenging cases such as fluid and fog.

Conclusion. We presented Tensor4D, a new method for learning high-quality neural representation for dynamic scenes from sparse-view videos or even a monocular video. To capture the spatio-temporal information in a compact and memory-efficient manner, we propose a novel hierarchical tri-projection decomposition method that models a 4D tensor with nine 2D feature planes. With proper design of training losses and regularization, our method provides an efficient yet effective solution to model the radiance fields of dynamic scenes. We believe our work can inspire future research towards low-cost, portable and immersive telepresence systems.

Acknowledgements. This paper is supported by National Key R&D Program of China (2022YFF0902200), the NSFC project No.62125107 and No.61827805 and the Beijing Science and Technology Project (Z211100004021006).

References

- [1] 2
- [2] The looking glass holographic display. <https://lookingglassfactory.com/>. 6
- [3] Naveed Ahmed, Christian Theobalt, Petar Dobrev, Hans-Peter Seidel, and Sebastian Thrun. Robust fusion of dynamic shape and normal capture for high-quality reconstruction of time-varying geometry. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. 2
- [4] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P.A. Lensch. Nerd: Neural reflectance decomposition from image collections. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [5] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. *arXiv preprint arXiv:2301.09632*, 2023. 3, 4
- [6] Eric Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *arXiv*, 2020. 2
- [7] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022. 1, 2, 3
- [8] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *Proceedings of the European Conference on Computer Vision*, 2022. 3
- [9] 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. 34(4), jul 2015. 1
- [10] Mingsong Dou, Philip Davidson, Sean Ryan Fanello, Sameh Khamis, Adarsh Kowdle, Christoph Rhemann, Vladimir Tankovich, and Shahram Izadi. Motion2fusion: Real-time volumetric performance capture. *ACM Trans. Graph.*, 36(6), nov 2017. 2
- [11] Yilun Du, Yinan Zhang, Hong-Xing Yu, Joshua B. Tenenbaum, and Jiajun Wu. Neural radiance flow for 4d view synthesis and video processing. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 14304–14314, 2021. 1, 2
- [12] Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiapeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. In *arXiv preprint arXiv:2205.15285*, 2022. 3, 6, 7, 8
- [13] 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. 3, 4
- [14] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5491–5500, 2022. 2
- [15] Wanshui Gan, Hongbin Xu, Yi Huang, Shifeng Chen, and Naoto Yokoya. V4d: Voxel for 4d novel view synthesis. In *arXiv preprint arXiv:2205.14332*, 2022. 3
- [16] Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular dynamic view synthesis: A reality check. In *In Neural Information Processing Systems (NeurIPS)*, 2022. 2
- [17] Paulo FU Gotardo, Tomas Simon, Yaser Sheikh, and Iain Matthews. Photogeometric scene flow for high-detail dynamic 3d reconstruction. In *Proceedings of the IEEE international conference on computer vision*, pages 846–854, 2015. 2
- [18] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d aware generator for high-resolution image synthesis. In *International Conference on Learning Representations*, 2022. 2
- [19] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, Danhang Tang, Anastasia Tkach, Adarsh Kowdle, Emily Cooper, Mingsong Dou, Sean Fanello, Graham Fyffe, Christoph Rhemann, Jonathan Taylor, Paul Debevec, and Shahram Izadi. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Trans. Graph.*, 38(6), nov 2019. 1
- [20] Yannan He, Anqi Pang, Xin Chen, Han Liang, Minye Wu, Yuexin Ma, and Lan Xu. Dynamic shape capture using multi-view photometric stereo. In *ACM SIGGRAPH Asia 2009 Papers*, New York, NY, USA, 2009. Association for Computing Machinery. 2
- [21] Hankyu Jang and Daeyoung Kim. D-tensorf: Tensorial radiance fields for dynamic scenes. *arXiv preprint arXiv:2212.02375*, 2022. 3, 4
- [22] Wobong Jang and Lourdes Agapito. Codenerf: Disentangled neural radiance fields for object categories. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12949–12958, 2021. 2
- [23] T. Kanade, P. Rander, and P.J. Narayanan. Virtualized reality: constructing virtual worlds from real scenes. *IEEE MultiMedia*, 4(1):34–47, 1997. 1
- [24] Hansung Kim, Jean-Yves Guillemaut, Takeshi Takai, Muhammad Sarim, and Adrian Hilton. Outdoor dynamic 3d scene reconstruction. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(11):1611–1622, 2012. 2
- [25] Jason Lawrence, Danb Goldman, Supreeth Achar, Gregory Major Blascovich, Joseph G. Desloge, Tommy Fortes, Eric M. Gomez, Sascha Häberling, Hugues Hoppe, Andy Huibers, Claude Knaus, Brian Kuschak, Ricardo Martin-Brualla, Harris Nover, Andrew Ian Russell, Steven M. Seitz, and Kevin Tong. Project starline: A high-fidelity telepresence system. *ACM Trans. Graph.*, 40(6), dec 2021. 2
- [26] Hao Li, Linjie Luo, Daniel Vlasic, Pieter Peers, Jovan Popović, Mark Pauly, and Szymon Rusinkiewicz. Temporally coherent completion of dynamic shapes. *ACM Trans. Graph.*, 31(1), feb 2012. 1
- [27] Tianye Li, Mira Slavcheva, Michael Zollhofer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, and Zhaoyang Lv. Neural 3d video synthesis from multi-view

- video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5521–5531, 2022. 2
- [28] 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. In *arXiv preprint, arXiv:2205.15723*, 2022. 1, 2, 3
- [29] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *Advances in Neural Information Processing Systems*, 2020. 2
- [30] Zhaoyang Lv, Kihwan Kim, Alejandro Troccoli, Deqing Sun, James M. Rehg, and Jan Kautz. Learning rigidity in dynamic scenes with a moving camera for 3d motion field estimation. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 484–501, Cham, 2018. Springer International Publishing. 2
- [31] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *CVPR*, 2021. 2
- [32] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2, 6
- [33] Thomas Muller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. 41(4), jul 2022. 3
- [34] Richard A. Newcombe, Dieter Fox, and Steven M. Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 343–352, 2015. 2
- [35] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [36] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [37] 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 International Conference on Computer Vision*, pages 5865–5874, 2021. 1, 2
- [38] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A higherdimensional representation for topologically varying neural radiance fields. *ACM Trans. Graph.*, 40(6), 2021. 2
- [39] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 2
- [40] Albert Pumarola, Enric Corona, and Francesc Moreno-Noguer Gerard Pons-Moll. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 1, 2, 3, 6, 7, 8
- [41] René Ranftl, Vibhav Vineet, Qifeng Chen, and Vladlen Koltun. Dense monocular depth estimation in complex dynamic scenes. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4058–4066, 2016. 2
- [42] Chris Russell, Rui Yu, and Lourdes Agapito. Video pop-up: Monocular 3d reconstruction of dynamic scenes. In *Computer Vision – ECCV 2014*, pages 583–598, Cham, 2014. Springer International Publishing. 2
- [43] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20154–20166. Curran Associates, Inc., 2020. 2
- [44] Ruizhi Shao, Liliang Chen, Zerong Zheng, Hongwen Zhang, Yuxiang Zhang, Han Huang, Yandong Guo, and Yebin Liu. Floren: Real-time high-quality human performance rendering via appearance flow using sparse rgb cameras. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–10, 2022. 2
- [45] Ruizhi Shao, Hongwen Zhang, He Zhang, Mingjia Chen, Yanpei Cao, Tao Yu, and Yebin Liu. Doublefield: Bridging the neural surface and radiance fields for high-fidelity human reconstruction and rendering. In *CVPR*, 2022. 2
- [46] Ruizhi Shao, Zerong Zheng, Hongwen Zhang, Jingxiang Sun, and Yebin Liu. Diffustereo: High quality human reconstruction via diffusion-based stereo using sparse cameras. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, pages 702–720. Springer, 2022. 2
- [47] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, 2019. 2
- [48] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *CVPR*, 2021. 2
- [49] Jonathan Starck and Adrian Hilton. Surface capture for performance-based animation. *IEEE Computer Graphics and Applications*, 27(3):21–31, 2007. 1
- [50] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. 2
- [51] Guoxing Sun, Xin Chen, Yizhang Chen, Anqi Pang, Pei Lin, Yuheng Jiang, Lan Xu, Jingya Wang, and Jingyi Yu. Neural free-viewpoint performance rendering under complex human-object interactions. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2021. 2
- [52] Jiaming Sun, Xi Chen, Qianqian Wang, Zhengqi Li, Hadar Averbuch-Elor, Xiaowei Zhou, and Noah Snavely. Neural

- 3D reconstruction in the wild. In *SIGGRAPH Conference Proceedings*, 2022. 2
- [53] Jingxiang Sun, Xuan Wang, Yong Zhang, Xiaoyu Li, Qi Zhang, Yebin Liu, and Jue Wang. Fenerf: Face editing in neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7672–7682, 2022. 2
- [54] Aparna Taneja, Luca Ballan, and Marc Pollefeys. Modeling dynamic scenes recorded with freely moving cameras. *Computer Vision – ACCV 2010*, page 613–626, 2011. 2
- [55] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhofer, 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 International Conference on Computer Vision*, pages 12959–12970, 2021. 1, 2
- [56] Chaoyang Wang, Ben Eckart, Simon Lucey, and Orazio Gallo. Neural trajectory fields for dynamic novel view synthesis. In *arXiv preprint, arXiv:2105.05994*, 2021. 1, 2, 3
- [57] Liao Wang, Ziyu Wang, Pei Lin, Yuheng Jiang, Xin Suo, Minye Wu, Lan Xu, and Jingyi Yu. *IButter: Neural Interactive Bullet Time Generator for Human Free-Viewpoint Rendering*, page 4641–4650. Association for Computing Machinery, New York, NY, USA, 2021. 2
- [58] 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 Conference on Computer Vision and Pattern Recognition*, pages 13524–13534, 2022. 3
- [59] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. 2, 5, 6, 7, 8
- [60] 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 (CVPR)*, pages 16210–16220, June 2022. 2
- [61] Liwen Wu, Jae Yong Lee, Yu-Xiong Wang Anand Bhattad, and David Forsyth. Diver: Real-time and accurate neural radiance fields with deterministic integration for volume rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 16200–16209, 2022. 3
- [62] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9421–9431, 2021. 1, 2, 3
- [63] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. 2
- [64] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33, 2020. 2
- [65] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5752–5761, 2021. 1, 2, 3
- [66] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgbd sensors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5746–5756, 2021. 2
- [67] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul Debevec, William T. Freeman, and Jonathan T. Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graph.*, 40(6), dec 2021. 2
- [68] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. Humannerf: Efficiently generated human radiance field from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7743–7753, June 2022. 2
- [69] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15893–15903, 2022. 2
- [70] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 286–301. Springer, 2016. 2