

Neural Parametric Gaussians for Monocular Non-Rigid Object Reconstruction

Devikalyan Das¹ Christopher Wewer² Raza Yunus^{1,2} Eddy Ilg¹ Jan Eric Lenssen²

¹Saarland University, Saarland Informatics Campus, Germany ²Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

{ddas, jlenssen}@mpi-inf.mpg.de



a) Novel View Synthesis on the D-NeRF Dataset

b) Comparison on an Unbiased4D Sequence.

Figure 1. We present **Neural Parametric Gaussians** (**NPGs**), a method for monocular non-rigid reconstruction of objects. **a**) Our method enables to produce high quality reconstructions in easy settings like the object-level D-NeRF scenes while **b**) also being able to handle challenging monocular settings much better than previous work through strong parametric low-rank regularization.

Abstract

Reconstructing dynamic objects from monocular videos is a severely underconstrained and challenging problem, and recent work has approached it in various directions. However, owing to the ill-posed nature of this problem, there has been no solution that can provide consistent, highquality novel views from camera positions that are significantly different from the training views. In this work, we introduce Neural Parametric Gaussians (NPGs) to take on this challenge by imposing a two-stage approach: first, we fit a low-rank neural deformation model, which then is used as regularization for non-rigid reconstruction in the second stage. The first stage learns the object's deformations such that it preserves consistency in novel views. The second stage obtains high reconstruction quality by optimizing 3D Gaussians that are driven by the coarse model. To this end, we introduce a local 3D Gaussian representation, where temporally shared Gaussians are anchored in and deformed by local oriented volumes. The resulting combined model can be rendered as radiance fields, resulting in high-quality photo-realistic reconstructions of the non-rigidly deforming objects. We demonstrate that NPGs achieve superior results compared to previous works, especially in challenging scenarios with few multi-view cues.¹

1. Introduction

Reconstructing 3D objects from 2D observations is a core problem in computer vision with numerous applications in several industries, such as the movie and game industry, AR/VR, and robotics. Tremendous progress has been seen in static scene reconstruction during the last few years. The real world is, however, dynamic, and most recorded scenes are captured in a casual setting, with sparse coverage from a single camera. Thus, addressing these two aspects during reconstruction is of fundamental importance.

The success of neural approaches on static scenes has encouraged their use for dynamic scene reconstruction from monocular videos, both in its classical [24, 40, 45] and hybrid [4, 10, 28] forms. These methods either learn a per-frame scene representation with limited time consistency [10, 24] or utilize a time-invariant canonical space, which is used to track the observations at each timestep [41, 45]. However, as pointed out in Gao et al. [12], they are evaluated on datasets that contain multi-view signals, such as camera teleportation-i.e., alternating samples from different cameras to construct a temporal sequence-and limited object motion, and their performance suffers drastically when evaluated on more realistic monocular sequences. Such sequences usually contain faster object motion compared to camera motion. Strong regularization is required in order to propagate information between different timesteps with the correct data association. As we will demonstrate

¹Project Page: https://geometric-rl.mpi-inf.mpg.de/npg/

on realistic sequences, current methods fail to provide such regularization. Traditionally, a possible solution to this has been to use geometry proxies in different forms [15], such as the SMPL [30] body model. However, these approaches require an appropriate model with a priori knowledge.

In this paper, we tackle the problem of template-free dynamic object reconstruction from realistic monocular captures. We propose Neural Parametric Gaussians (NPGs) as a two-step approach, which first learns the coarse, deforming geometry and then uses it as a proxy for reconstruction. In the first step, we obtain a coarse neural parametric model for the observed object at each timestep, generated using a low-rank decomposition of deformation and a set of point basis vectors. The basis vectors force the model to share information between timesteps, thus providing regularization for the sparsely observed dynamic regions. In the second step, we represent fine-level geometry and appearance details by initializing and optimizing 3D Gaussians on top of the coarse point template, inspired by the recent breakthrough provided by the 3D Gaussian splatting approach for static scenes [18]. For handling dynamic objects, we define Gaussians in oriented local volumes such that they can be driven by the coarse deformation model from the first stage. In short, our contributions are:

- A two-stage reconstruction approach that learns coarse deformation in stage one and uses it as a constraint for reconstruction in stage two.
- A coarse parametric point model based on a low-rank deformation that provides regularization and correspondences over time for object reconstruction.
- A 3D Gaussian splatting approach, where Gaussians live in deforming local volumes, generalizing 3D Gaussians from static to dynamic scenes efficiently.
- We show that our model improves on the state of the art of non-rigid novel view synthesis from a monocular camera, especially in challenging cases with few multi-view cues.

2. Related Work

We look at a brief overview of the recent advances in nonrigid reconstruction from the perspective of scene representations, deformation modeling, and supervision, focusing on monocular and object-level reconstruction methods. For a more in-depth discussion, please refer to the recent survey by Yunus *et al.* [74].

2.1. Neural Rendering of Dynamic Scenes

NeRF [32] has been extended to model dynamic scenes both using a global MLP representation and as a hybrid representation, where neural features are stored at the nodes of a discrete data structure. These representations capture the scene dynamics in one of two ways. One approach, dubbed *Space-Time Neural Fields*, directly adds an extra time dimension to the scene representation to reconstruct the dynamic scene from multi-view [1, 21, 23, 26, 42, 50, 59-61, 63, 70] or monocular [4, 8, 10, 11, 24, 25, 48, 51, 67] video. Hybrid methods further achieve acceleration by parameterizing the 4D scene representation using voxel grids [21, 42, 51, 59–61, 63] or planar factorization [1, 4, 10, 26, 48, 70]. Such methods often rely on estimated depth maps and pre-computed optical flow for local motion modeling but fail to propagate information globally, required to render novel views from viewpoints significantly different than the observations, making them suitable mostly for forward-facing videos. Similar to ours, Hexplane [4] shares information between timesteps by proposing a low-rank temporal basis for the spatially decomposed volumes. However, they optimize the basis together with the fine-level geometry and appearance while our two-stage approach provides stronger regularization by first obtaining a coarse template and then optimizing fine details on top.

The other paradigm, called Deformable NeRFs, introduces a 4D deformation field that maps the observations at each timestep to a canonical space, i.e. backward warping, ensuring temporal correspondences in contrast to the previous paradigm but restricting large deviations and topological changes from the canonical field [7, 40, 41, 45, 55, 58]. A few approaches represent the 4D deformation and the canonical radiance field with a voxel-based hybrid representation for fast multi-view reconstruction [28, 56]. TiNeu-Vox [9] further accelerates monocular reconstruction by using a very light MLP for the backward deformation field, compensating for compression by enhancing the scene representation through temporal embeddings. Backward warping is not smooth; hence, it has difficulties generalizing to sparse capture settings for novel view synthesis, where strong regularization is required. Recently, ForwardFlowD-NeRF [13] proposed the use of time-dependent voxel features in the canonical space for learning a forward deformation field to each timestep instead of the other way around, leading the MLP to model motion more smoothly. In contrast, our optimized coarse model provides explicit regularization for reconstruction in the second stage.

2.2. Object-level Non-Rigid Reconstruction

Object-level approaches employ masks to separate the dynamic object of interest from the background and mostly focus on surface modeling. Using classical representations, Shape-from-Template methods [5, 17] rely on preacquired full object templates, while template-free online methods [2, 6, 27, 34] use RGB-D input to reconstruct the object's geometry. A few recent neural rendering approaches focus on surface modeling of masked objects from monocular RGB [15] and RGB-D [3] videos by replacing the density-based representation with SDF modeling from NeuS [62]. Unbiased4D [15] also utilizes a mesh proxy to model large deformations. A couple of methods [49, 53] employ data-driven category-level priors for shape and appearance, along with articulation [53] or scene flow [49] prediction, whereas our template prediction does not rely on any learned data priors but is purely optimization-based.

Neural Parametric Models utilize an auto-decoded MLP to learn the articulation space-either represented by scene flow [33, 38, 39] or volumetric bone transformations as in BANMO [71]-for a given shape from monocular videos. Similar to ours, KeypointTransporter [36] reconstructs coarse 3D point clouds for objects, parameterized by a low-rank deformation basis. In contrast to such methods, which focus on coarse geometry only, our approach focuses on high-quality view synthesis using a deformable parametric model. A few approaches utilize either skeletons [22, 29, 35, 73] or surface templates [14, 43, 68, 77] from parametric models like SMPL [30] to guide the motion field and learn the canonical radiance field on top of such templates. While the template provides motion regularization, these approaches require a priori information about what is to be reconstructed, which does not generalize to sequences in the wild. Our coarse point template, on the other hand, is derived solely from the observations.

2.3. Point-based Representation and Rendering

A few approaches utilize point-based neural rendering, introduced for static scenes by Point-NeRF [69], for dynamic avatar modeling from multi-view [57] and monocular [52] videos. The recent trend has been shifting towards using purely point-based explicit representations for their efficiency, showing great results from multi-view input for dynamic surface modeling [44] and view synthesis using point-based rasterization [75]. 3D-GS [18], a recent breakthrough approach for static scenes, models the scene with numerous volumetric 3D Gaussians and proposes a Gaussian splatting-based rasterization technique to deliver stateof-the-art results in terms of view synthesis quality and rendering speed. Our approach extends the static 3D-GS method to dynamic objects.

Concurrent Works. In terms of modeling dynamic scenes with 3D Gaussians, several concurrent works have appeared on arXiv.org. Dynamic3DGS [31] utilizes dense multi-view inputs to track the rotation and position of 3D Gaussians initialized in the first frame across timesteps. Two similar works, Deformable3DGS [72] and 4DGaussianSplatting [66] introduce a canonical space where 3D Gaussians are initialized while using a neural deformation field to track the position, rotation, and scale of each Gaussian across timesteps. 4DGaussianSplatting additionally represents the deformation field using a decomposed Hexplane representation. In contrast, our approach introduces the 3D Gaussian scene representation specifically to the underconstrained setting of sparse monocular captures, by initializing the Gaussians from the underlying coarse point

template and using the deformation of the template as a guide for optimization of the 3D Gaussians. Also, in contrast to some works [66, 72], we do not apply MLPs to all Gaussians in a scene, which leads to better scalability.

3. Method

The presented method is a two-stage approach. First, we obtain constraints for reconstruction by fitting a neural parametric point model to monocular video sequences (c.f. Sec. 3.1). Then, the second stage uses the learned constraints from stage 1 to solve the underconstrained task of monocular, non-rigid reconstruction (c.f. Sec. 3.2). Specifically, we formulate 3D Gaussians as a function of the underlying point template, optimizing the parameters with respect to local subspaces.

3.1. Coarse Parametric Point Model

The goal of the first stage is to obtain a coarse, non-rigid point model that follows the movement of the object, given a sequence of monocular input views $\{\mathbf{I}_i\}_{i=1}^N$ with masks $\{\mathbf{M}_i\}_{i=1}^N$ and camera poses. For real videos the poses are obtained via COLMAP [46, 47]. We do not make any assumptions of 3D location of the object, but we require the object to stay in view. Note that, given recent advances in segmentation, in contrast to previous work, we consider the mask to be a weak assumption and demonstrate that our method works with masks obtained from Segment Anything (SAM) [20] in the supplemental materials.

3.1.1 Representation

We represent our coarse, non-rigid model as a set of points $\mathbf{P}^t \in \mathbb{R}^{M \times 3}$ for $t \in \{1, ..., N\}$, which are the output of a low-rank deformation basis

$$\mathbf{P}^{t} = \sum_{k=1}^{K} \alpha_{k}^{t} \mathbf{B}_{k}, \tag{1}$$

where $\{\mathbf{B}_k \in \mathbb{R}^{M \times 3}\}_{k=1}^K$ is a learnable point basis and $\alpha_k^t \in \mathbb{R}$ are low-rank coefficients. The latter is produced by an MLP $f_{\theta} : \mathbb{R}^d \to \mathbb{R}^K$, mapping the temporally encoded time $\phi(t) \in \mathbb{R}^d$ to coefficients over time, following the ideas of multiple recent works [36, 37, 65]. The basis is randomly initialized. This representation naturally provides correspondences over time, as the point cloud is smoothly deformed between frames. By choosing K, we have control over the rigidity of the model. Intuitively, each basis vector can represent the object in a certain pose, and the linear combination of these shapes defines the extent of deformations that can be modeled, thus providing regularization. We choose $K \ll T$ and evaluate different choices in Sec. 4.2. Additionally, the points are equipped with coarse colors $\mathbf{C} \in \mathbb{R}^{N \times 3}$ and random features $\mathbf{E} \in \mathbb{R}^{N \times d}$ for identification purposes.



Figure 2. Overview of our method. We present a two-stage method. In stage 1 (left) we learn a coarse point model, which is parameterized through low-rank coefficients from an MLP. In stage 2 (right), we optimize 3D Gaussians in local volumes, defined by the point sets. The figure distinguishes between parts that are shared over time (\blacksquare), individual for each time step (\blacksquare), and fixed-function (\blacksquare). MLP weights θ , Gaussian interpolation weights w, scales S, rotations R and harmonic coefficients h are shared over time and the deformation is purely modeled by the low-rank coefficients α_i , leading to a different coarse point model for each frame.

3.1.2 Optimization

.

Given a sequence of input views $\{\mathbf{I}_i\}_{i=1}^N$ and masks $\{\mathbf{M}_i\}_{i=1}^N$, we follow a similar procedure as in KeyTr [36] and find optimal MLP weights $\hat{\theta}$, a deformation basis $\{\hat{\mathbf{B}}_k\}_{k=1}^K$ and colors $\hat{\mathbf{C}}$ as

$$\hat{\theta}, \{\hat{\mathbf{B}}_k\}, \hat{\mathbf{C}}^t = \operatorname*{arg\,min}_{\theta, \{\mathbf{B}_k\}, \mathbf{C}^t} (\mathcal{L}_{\mathrm{M}} + \mathcal{L}_{\mathrm{R}} + \mathcal{L}_{\mathrm{OF}}), \quad (2)$$

with the individual losses being described in the following. For the loss weights, please refer to the supplementals.

Mask Loss. For each frame *i*, we sample points \mathbf{P}_{M}^{i} uniformly from the given mask and minimize the Chamfer distance $\mathcal{L}_{M} = CD(\mathbf{P}_{M}^{i}, \pi(\mathbf{P}^{i}))$ to the point model projected to the image plane by π . We found this to be more efficient and equally good as a Sinkhorn matching loss [36].

Rigidity Loss. We add a rigidity loss that keeps neighborhood point distances of a frame t similar to that of a reference frame r:

$$\mathcal{L}_{\mathbf{R}} = \sum_{i=2}^{M} \sum_{j \in \mathcal{N}(i)} \| \| \mathbf{P}_{i}^{r} - \mathbf{P}_{j}^{r} \|_{2} - \| \mathbf{P}_{i}^{t} - \mathbf{P}_{j}^{t} \|_{2} \|_{2}^{2}, \quad (3)$$

for all $t \in \{1, ..., N\}$, where $\mathcal{N}(i)$ denotes the neighborhood points of point *i*, given a radius criterion, and *r* is a reference frame selected manually from the dataset.

Optical Flow Loss. We use RAFT [54] as an off-the-shelf optical flow estimator to align point movement between frames via estimated optical flow, as done by KeyTr [36]. The flow consistency loss between two frames is given as

$$\mathcal{L}_{\mathrm{OF}}^{t \mapsto t+1} = \|\mathbf{M}^t \odot [R(\mathbf{P}^t, \mathbf{E}) - R^B(\mathbf{P}^{t+1}, \mathbf{E})])\|_{\epsilon} \quad (4)$$

where $R(\mathbf{P}^t, \mathbf{E})$ renders the point cloud \mathbf{P}^t with random descriptors \mathbf{E} , and $R^B(\mathbf{P}^{t+1}, \mathbf{E})$ renders, applies the backward optical flow and bilinear samples the result. $||x||_{\epsilon}$ is an element-wise Huber loss with threshold $\epsilon = 0.01$.

3.2. Neural Parametric Gaussians

Given our parametric point model, we perform a detailed reconstruction of the monocular sequence by letting the coarse model drive 3D Gaussians that live in oriented local volumes defined by point sets. We first define our local volumes with Gaussians in Sec. 3.2.1. Then, the rendering process is described in Sec. 3.2.2, the densification process in Sec. 3.2.3, and the optimization process in Sec. 3.2.4.

3.2.1 Representation

When defining the local volumes, the goal is to make as few assumptions about the object as possible and, particularly, to avoid the assumption of a surface. Thus, we consider local unstructured point sets as volumes.

Local Volumes. We define an *oriented local volume* as $(\mathbf{V}(i), \mathbf{T}(i))$, where $\mathbf{V}(i) \in \mathbb{R}^{k \times 3}$ is the point with index *i* plus its k - 1 nearest neighbors and $\mathbf{T}(i) \in SO(3)$ is the local volume orientation. We choose k = 20 in our experiments and the local orientation $\mathbf{T}(i)$ heuristically by taking the orientation of the triangle of point *i* and its closest 2 neighbors. The neighbors defining the triangle are fixed over time (please refer to the supplemental materials for a detailed description). Then, the full set of local volumes at time *t* is given as

$$\mathcal{V} = \{ (\mathbf{V}^t(i), \mathbf{T}^t(i)) \mid \forall i \in \{1, ..., M\} \}.$$
(5)

The set of volumes is fixed over time, however, their point positions $\mathbf{V}^t(i)$ change from frame to frame. Note that the

absolute orientation of the local volumes is not relevant, as the rotations \mathbf{R} of individual Gaussians can adjust accordingly (see next paragraph). Importantly, the $\mathbf{T}(i)$ are equivariant to local deformation: if the local point set is rotated, the rotation-defining triangle rotates, resulting in $\mathbf{T}(i)$ changing by the same rotation, as desired.

3D Gaussians. Each local volume can store an arbitrary number of 3D Gaussians [18], which we denote as \mathcal{G}_i . Having a different number of Gaussians for each volume allows the model to adapt to the amount of high-frequency details in individual object parts. A single Gaussian $g \in \mathcal{G}_i$ is defined as tuple $g = (\mathbf{w}, \mathbf{S}, \mathbf{R}, \mathbf{h})$, where \mathbf{S}, \mathbf{R} , and \mathbf{h} are scale matrix, rotation matrix and spherical harmonic coefficients as in 3D Gaussian splatting [18], respectively. Instead of representing the 3D Gaussian position in world space, we represent its position as a set of barycentric interpolation weights $\mathbf{w} \in \mathbb{R}^k$ applied to the neighboring points $\mathbf{V}(i)$. Before rendering a Gaussian belonging to volume i, we find the final Gaussian position \mathbf{x}^t and rotations \mathbf{R}^t at time t:

$$\mathbf{x}^{t} = \sum_{j=1}^{k} \operatorname{Softmax}(\mathbf{w})_{j} \cdot \mathbf{V}^{t}(i)_{j}, \text{ and } \bar{\mathbf{R}}^{t} = \mathbf{T}^{t}(i) \cdot \mathbf{R},$$
 (6)

via a linear combination of volume-defining points and rotation into the local coordinate frame, respectively. The process can be understood as a simplified variant of cage-based deformation transfer [16]. It is important to note that Gaussian parameters are shared across all t and that temporal changes are modeled purely by the coarse point set.

3.2.2 Rendering

The sets of Gaussians from all volumes in \mathcal{V} are rendered using the differentiable rasterizer provided by Kerbl *et al.* [18]. It is important to note that, in contrast to concurrent works [66, 72], our model does not apply networks to Gaussians. Since the number of Gaussians within one scene can become very large, this avoids heavy runtime increases for larger scenes. We only have a network driving our initial point template, which has a comparably low number of points, while Gaussians just follow their local movement and have directly optimizable parameters that are shared over the full sequence. Further, given an optimized representation, it is straightforward to extract all Gaussians and store them for real-time rendering.

3.2.3 Gaussian Initialization and Densification

When initializing our model, we fill each volume with a small number of Gaussians. During training, we apply densification and pruning procedures from Kerbl *et al.* [18]. Importantly, when splitting a Gaussian, we assign new Gaussians to the same volume as the original Gaussian and apply a tiny bit of noise to weights w of both. In practice,

we store all Gaussians in a global list and keep a volume index for each Gaussian, allowing different amounts of Gaussians for each volume.

3.2.4 Optimization

In stage 2, we mainly optimize the individual Gaussians. Thus, we find optimal weights $\hat{\mathbf{w}}$, rotations $\hat{\mathbf{R}}$, scales $\hat{\mathbf{S}}$ and spherical harmonic coefficients $\hat{\mathbf{h}}$ as

$$\hat{\mathbf{w}}, \hat{\mathbf{R}}, \hat{\mathbf{S}}, \hat{\mathbf{h}} = \operatorname*{arg\,min}_{\mathbf{w}, \mathbf{R}, \mathbf{S}, \mathbf{h}} (1 - \lambda) \mathcal{L}_1 + \mathcal{L}_{\text{D-SSIM}},$$
 (7)

where the loss functions are *l*1-distance and a structural similarity loss, as used by Kerbl *et al.* [18]. Optionally, we can finetune all parameters from stage 1 based on these rendering losses, including a weak regularization to prevent large changes. While this is not strictly necessary, it slightly improves the details of the resulting reconstruction.

4. Experiments

In this section, we evaluate our method on synthetic and real monocular datasets and perform ablation studies to demonstrate its effectiveness in reconstructing realistically captured dynamic scenes.

Datasets. First, we evaluate our method on the commonly used D-NeRF benchmark dataset [45] and show that it reaches state-of-the-art performance on objects which fit our setting. Note that the dataset contains unrealistic multiview cues, like camera teleportation, which aid the reconstruction. After that, we provide quantitative, qualitative, and ablation studies on the more realistically captured Unbiased4D dataset [15] to demonstrate that our method provides adequate regularization for this underconstrained setting, where a camera is moving slowly around an object. We use the Effective Multi-view Factor (EMF) metric from Gao *et al.* [12] to determine the amount of multi-view cues present in datasets. D-NeRF has an EMF ω value of 2135.45, while the average EMF ω of Unbiased4D is 24.93.

Comparison Methods. Existing object-level methods mostly focus on geometry reconstruction and neglect appearance. So, we choose state-of-the-art methods for monocular scene-level reconstruction and use object masks for comparison. We mainly compare against Hexplane [4], which uses planar decomposition and information sharing for the 4D volume, and TiNeuVox-B [9], which utilizes a voxel-parametrized hybrid field for deformation modeling. We use the base version (B) of TiNeuVox, which is a bit slower but more accurate than the small version. Further comparisons are made against HyperNeRF [41], Nerfies [40] and D-NeRF [45] to put our performance more in context. We report the PSNR, SSIM [64], and LPIPS [76] metrics for all experiments to quantitatively evaluate the view synthesis quality.

	Jumping Jacks		Hell Warrior			Hook			Stand Up			
Method	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR ↑	SSIM↑	LPIPS↓
D-NeRF [45]	32.80	0.98	0.03	25.02	0.95	0.06	29.25	0.96	0.11	32.79	0.98	0.02
TiNeuVox [9]	34.23	0.98	0.03	28.17	0.97	0.07	31.45	0.97	0.05	35.43	0.99	0.02
HexPlane [4]	31.65	0.97	0.04	24.24	0.94	0.07	28.71	0.96	0.05	34.36	0.98	0.02
Ours	34.06	0.99	0.03	38.73	0.98	0.04	33.73	0.98	0.03	37.95	0.99	0.02
	1			1								
		Mutant			T-Rex			Lego			Average	
Method	PSNR↑	Mutant SSIM↑	LPIPS↓	PSNR↑	T-Rex SSIM↑	LPIPS↓	PSNR↑	Lego SSIM↑	LPIPS↓	PSNR↑	Average SSIM↑	LPIPS↓
Method D-NeRF [45]	PSNR↑ 31.29	Mutant SSIM↑ 0.97	LPIPS↓ 0.02	PSNR↑ 31.75	T-Rex SSIM↑ 0.97	LPIPS↓ 0.03	PSNR↑ 21.64	Lego SSIM↑ 0.83	LPIPS↓ 0.16	PSNR↑ 29.22	Average SSIM↑ 0.95	LPIPS↓ 0.06
Method D-NeRF [45] TiNeuVox [9]	PSNR↑ 31.29 33.61	Mutant SSIM↑ 0.97 0.98	LPIPS↓ 0.02 0.03	PSNR↑ 31.75 32.70	T-Rex SSIM↑ 0.97 0.98	LPIPS↓ 0.03 0.03	PSNR↑ 21.64 25.02	Lego SSIM↑ 0.83 0.92	LPIPS↓ 0.16 0.07	PSNR↑ 29.22 31.52	Average SSIM↑ 0.95 0.97	LPIPS↓ 0.06 0.04
Method D-NeRF [45] TiNeuVox [9] HexPlane [4]	PSNR↑ 31.29 33.61 33.79	Mutant SSIM↑ 0.97 0.98 0.98	LPIPS↓ 0.02 0.03 0.03	PSNR↑ 31.75 32.70 30.67	T-Rex SSIM↑ 0.97 0.98 0.98	LPIPS↓ 0.03 0.03 0.03	PSNR↑ 21.64 25.02 25.22	Lego SSIM↑ 0.83 0.92 0.94	LPIPS↓ 0.16 0.07 0.04	PSNR↑ 29.22 31.52 29.81	Average SSIM↑ 0.95 0.97 0.96	LPIPS↓ 0.06 0.04 0.04

Table 1. **Novel view synthesis on the synthetic D-NeRF dataset.** We evaluate our method quantitatively on scenes from the synthetic D-NeRF dataset, which fit our object-centric setting. It can be seen that our NPGs reach state-of-the-art performance. Even if the differences in metrics are small, the qualitatively results in Fig. 3 show clear differences in level of detail and correspondences. The Lego and T-Rex sequences pose a special challenge to our object-level method, as they contain a large static ground plane. We obtain results on these sequences by obtaining masks for dynamic objects only via Segment-Anything [20], as detailed in the supplementary materials.

Implementation Details. For stage one optimization, we use a learning rate of 0.0005 with a cosine annealing scheduler and warm-up, a batch size of 10, the Adam optimizer [19] and optimize for a total of 100k iterations. We employ 1500 points to model the coarse object geometry and a 6-layer MLP to obtain the coefficients for the deformation basis. In stage two, we mainly adopt the learning rate schedule from 3DGS [18]. The LR for interpolation weights w is the same as the default learning rate for point position as used in 3DGS, while we adapt the learning rate of the Gaussian scale to 0.0005 to match the smaller scenes. We train for a total of 70k iterations and use an exponential learning rate scheduler for the Gaussian scales, with a very high rate initially and a gradual decay to 0.0005 in the end. For finetuning stage 1, we use a learning rate of 0.0001.

4.1. Results and Comparison

Comparison on D-NeRF Dataset. The D-NeRF dataset [45] is a monocular dataset with 360 views of synthetic objects. It provides eight sequences with 200 training views and 20 test views. The sequences contain teleporting camera motion, which provides multi-view cues for the reconstruction, relaxing the difficulty for the dataset [12]. We compare renderings with resolution 400×400 as done by previous work and provide results on the full resolution of 800×800 in the supplemental materials. Although D-NeRF is not our target setting because of strong multiview cues, it provides verification that our method can reconstruct highly detailed objects from high-quality observations. We provide quantitative results for scenes that fit our object-level setting in Tab. 1. The results show that our method slightly outperforms the previous state-of-the-art approaches. A qualitative comparison is shown in Fig. 3 and the quality of our optimized templates can be seen in the bottom row on Fig. 4. Since strong regularization is not required for these sequences, most of the gains come

Method	PSNR ↑	SSIM↑	LPIPS↓
Nerfies [40]	17.36	0.87	0.13
HyperNeRF [41]	18.56	0.88	0.12
TiNeuVox [9]	15.92	0.848	0.155
HexPlane [4]	16.327	0.85	0.16
Ours	22.348	0.905	0.095

Table 2. Novel view synthesis on the Unbiased4D dataset. We clearly outperform previous methods on this challenging dataset, where the amount of multi-view cues is very low.

from improved details in representation. However, as seen in the second example, there are also some cases where our method prevents inconsistencies in geometry, where previous methods fail to do so. It should also be noted that HexPlane and TiNeuVox are capable of representing high-quality details in theory and that blurriness is a result of non-rigid deformation. This suggests that our method deals better with representing high-frequency details under deformation, even with stronger regularization.

Comparison on Unbiased4D Dataset. To evaluate realistically captured sequences that require strong regularization, we utilize the monocular dataset provided by Unbiased4D [15]. The original paper required geometry proxies to successfully reconstruct the sequences. It provides five sequences of both synthetic and real objects, each providing 150 training views. We use masks for us and all baselines to ensure a fair comparison. Only one of the sequences, the synthetic cactus, provides ground truth novel views for evaluating reconstruction quality quantitatively. We show the results in Tab. 2. It can be seen that we outperform all previous methods. The reasons for this can be seen in Fig. 6. While previous methods fail to capture the geometry correctly, NPGs keep it intact over the sequence and successfully fit details. The coarse point model learned in the first stage successfully serves as a proxy for reconstruction



Figure 3. Qualitative comparison on novel views of the D-NeRF dataset. We can see that our method produces more detailed reconstruction than previous work. Also, even with multi-view cues in D-NeRF, previous methods fail to always keep correct correspondences, as seen in the second example around the feet. In contrast, our NPGs keeps the shape coherent at all times and captures high frequency details under deformation.

in the second stage, resolving ambiguities. The smoothness of our motion modelling can be seen in Fig. 4. It is apparent that, in contrast to baselines, our model performs correct tracking, which is important for consistent reconstruction.

Rendering speed Due to the efficient Gaussian splatting rasterizer and no MLPs applied per Gaussian, NPGs have very competitive rendering times. On an NVidia A40, we render a frame in approximately 0.05s (20fps) in 400×400 image resolution. In comparison, TiNeuVox requires 2.15s



Figure 4. **Point trajectory visualization.** Our coarse parametric model automatically provides point trajectories, which in turn demonstrates the quality and smoothness of our optimized templates here. **Top Row:** Synthetic Cactus, Real Cactus and Synthetic Human sequences from the Unbiased4D dataset. **Bottom Row:** Jumping Jack, Stand Up and Hook sequences from the D-NeRF dataset. Note that the human on the top right is sliding with constant speed in this sequence, which is visible in the trajectories.



Figure 5. Rendered depth from optimized NPGs. We render depth maps from optimized models, showing consistent geometry.

Method	PSNR ↑	SSIM↑	LPIPS↓
Ours w/o Rigidity Loss	20.73	0.89	0.11
Ours w/o Template Fine Tuning	20.92	0.89	0.097
Ours ($K = 100$)	21.80	0.89	0.10
Ours ($K = 10$)	16.63	0.86	0.14
Ours $(K = 2)$	15.67	0.87	0.15
Ours ($K = 25$)	22.348	0.905	0.095

Table 3. Ablation study on the synthetic cactus sequence from Unbiased4D dataset. The first two rows show the analysis for coarse parametric model optimization while the next three rows demonstrate the effect of deformation basis size K.

per frame and Hexplane 0.94*s* per frame. Our rendering speed can be further increased by extracting the Gaussian representation explicitly and utilizing the real-time rendering of 3D Gaussian splatting [18].



Figure 6. **Qualitative comparison on the Unbiased4D dataset.** The Unbiased4D dataset is a challenging monocular video dataset where the amount of multi-view cues is very low. We show two different novel views for each scene. It can be seen that previous methods fail to correctly reconstruct the objects, while our NPGs keep the coarse geometry intact and provide novel views with high frequency details. On the most right, we show the input view for the given timeframe. Our method provides consistent reconstructions of far away novel views.



Figure 7. Qualitative ablation study on the Unbiased4D dataset. We show ablation on a novel view of the synthetic cactus sequence at a particular timestep. The regularization provided by our low-rank deformation basis can be clearly seen with respect to the ground truth novel view on the right.

4.2. Ablation Study

Coarse Parametric Model Optimization. We provide ablation for two aspects of our coarse parametric point model optimization quantitatively in Tab. 3 and qualitatively in Fig. 7. First, we show the effect of utilizing the rigidity loss during stage 1 optimization, which helps keep the template consistent across timesteps. Next, we demonstrate the effect of fine-tuning stage 1 parameters during stage 2 optimization, which, while not necessary, improves the accuracy of the template nonetheless.

Deformation Basis Rigidity. We compare the effect of different regularization strengths provided by the low-rank deformation basis in Tab. 3 and Fig. 7, by controlling the basis size K. A smaller basis size over-regularizes the reconstruction, restricting non-rigid deformations, while a larger basis makes the model overfit to the observations more, hindering novel view synthesis. We find the optimal size to be 25 for the synthetic cactus sequence.



Figure 8. Failure case of a flat template and wrong deformation.

Limitations. We tackle the highly ill-posed problem of reconstructing **non-rigid** objects from **monocular** videos. Thus, there are limitations depending on the complexity of sequences, e.g., w.r.t. camera movement, lighting, speed and extent of deformations. Fig. 8, for example, shows results where the template collapsed to a flat surface, which we attribute to the camera being very static.

5. Conclusion

We introduce Neural Parametric Gaussians, a high-fidelity view synthesis approach for objects captured with monocular camera. Our two-stage optimization obtains a coarse parametric point model based on a low-rank deformation basis, providing strong regularization for consistent novelview synthesis from sparse observations, while the use of 3D Gaussians in conjunction with the template enables modeling fine geometric and appearance details efficiently. NPG is the first approach that employs the regularization power of neural parametric models for high-quality novel view synthesis. We demonstrated that, in contrast to previous work in monocular non-rigid reconstruction, we can achieve consistent models in challenging settings.

Acknowledgments. R. Yunus was supported by the Max Planck & Amazon Science Hub. Thanks to all supporters.

References

- Benjamin Attal, Jia-Bin Huang, Christian Richardt, Michael Zollhoefer, Johannes Kopf, Matthew O'Toole, and Changil Kim. Hyperreel: High-fidelity 6-dof video with rayconditioned sampling. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16610–16620, 2023.
- [2] Aljaz Bozic, Pablo Palafox, Michael Zollhöfer, Angela Dai, Justus Thies, and Matthias Nießner. Neural non-rigid tracking. In Int. Conference on Neural Information Processing Systems (NeurIPS), pages 18727–18737, 2020. 2
- [3] Hongrui Cai, Wanquan Feng, Xuetao Feng, Yan Wang, and Juyong Zhang. Neural surface reconstruction of dynamic scenes with monocular rgb-d camera. In *Int. Conference* on Neural Information Processing Systems (NeurIPS), pages 967–981, 2022. 2
- [4] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 130–141, 2023. 1, 2, 5, 6
- [5] David Casillas-Perez, Daniel Pizarro, David Fuentes-Jimenez, Manuel Mazo, and Adrien Bartoli. The isowarp: the template-based visual geometry of isometric surfaces. pages 2194–2222, 2021. 2
- [6] Haonan Chang, Dhruv Metha Ramesh, Shijie Geng, Yuqiu Gan, and Abdeslam Boularias. Mono-star: Mono-camera scene-level tracking and reconstruction. *arXiv pre-print*, 2023. 2
- [7] Jaesung Choe, Christopher Choy, Jaesik Park, In So Kweon, and Anima Anandkumar. Spacetime surface regularization for neural dynamic scene reconstruction. In *Int. Conference* on Computer Vision (ICCV), pages 17871–17881, 2023. 2
- [8] Yilun Du, Yinan Zhang, Hong-Xing Yu, Joshua B Tenenbaum, and Jiajun Wu. Neural radiance flow for 4d view synthesis and video processing. In *Int. Conference on Computer Vision (ICCV)*, pages 14304–14314. IEEE Computer Society, 2021. 2
- [9] Jiemin Fang, Taoran Yi, Xinggang 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, 2022. 2, 5, 6
- [10] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 12479–12488, 2023. 1, 2
- [11] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *Int. Conference on Computer Vision (ICCV)*, 2021. 2
- [12] Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular dynamic view synthesis: A reality check. In *Int. Conference on Neural Information Processing Systems (NeurIPS)*, 2022. 1, 5, 6
- [13] Xiang Guo, Jiadai Sun, Yuchao Dai, Guanying Chen, Xiaoqing Ye, Xiao Tan, Errui Ding, Yumeng Zhang, and Jingdong

Wang. Forward flow for novel view synthesis of dynamic scenes. In *Int. Conference on Computer Vision (ICCV)*, pages 16022–16033, 2023. 2

- [14] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *Proceedings of the European conference on computer vision (ECCV)*, 2022. 3
- [15] Erik Johnson, Marc Habermann, Soshi Shimada, Vladislav Golyanik, and Christian Theobalt. Unbiased 4d: Monocular 4d reconstruction with a neural deformation model. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 6597–6606, 2023. 2, 5, 6
- [16] Tao Ju, Scott Schaefer, and Joe Warren. Mean value coordinates for closed triangular meshes. ACM Trans. Graph., 24 (3):561–566, 2005. 5
- [17] Navami Kairanda, Edith Tretschk, Mohamed Elgharib, Christian Theobalt, and Vladislav Golyanik. f-sft: Shapefrom-template with a physics-based deformation model. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 3948–3958, 2022. 2
- [18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics, 42 (4), 2023. 2, 3, 5, 6, 7
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014. 6
- [20] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. arXiv preprint arXiv:2304.02643, 2023. 3, 6
- [21] Lingzhi Li, Zhen Shen, Zhongshu Wang, Li Shen, and Ping Tan. Streaming radiance fields for 3d video synthesis. In Int. Conference on Neural Information Processing Systems (NeurIPS), pages 13485–13498, 2022. 2
- [22] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhöfer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. In *European Conhference on Computer Vision (ECCV)*, pages 419– 436. Springer, 2022. 3
- [23] 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 *Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 5521–5531, 2022. 2
- [24] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6498–6508, 2021. 1, 2
- [25] Zhengqi Li, Qianqian Wang, Forrester Cole, Richard Tucker, and Noah Snavely. Dynibar: Neural dynamic image-based rendering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4273–4284, 2023. 2
- [26] Haotong Lin, Sida Peng, Zhen Xu, Tao Xie, Xingyi He, Hujun Bao, and Xiaowei Zhou. High-fidelity and real-time novel view synthesis for dynamic scenes. In ACM SIG-GRAPH, 2023. 2

- [27] Wenbin Lin, Chengwei Zheng, Jun-Hai Yong, and Feng Xu. Occlusionfusion: Occlusion-aware motion estimation for real-time dynamic 3d reconstruction. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1736–1745, 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 *Int. Conference on Neural Information Processing Systems (NeurIPS)*, pages 36762– 36775, 2022. 1, 2
- [29] Jia-Wei Liu, Yan-Pei Cao, Tianyuan Yang, Eric Zhongcong Xu, Jussi Keppo, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Hosnerf: Dynamic human-object-scene neural radiance fields from a single video. arXiv pre-print, 2023. 3
- [30] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multiperson linear model. pages 248:1–248:16. ACM, 2015. 2, 3
- [31] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *Int. Conference on Computer Vision (ICCV) Workshop*, 2023. 3
- [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 *European Configuration Computer Vision* (ECCV), 2020. 2
- [33] Mirgahney Mohamed and Lourdes Agapito. Gnpm: Geometric-aware neural parametric models. In Int. Conference on 3D Vision, pages 166–175. IEEE, 2022. 3
- [34] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 343–352, 2015. 2
- [35] Atsuhiro Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Unsupervised learning of efficient geometry-aware neural articulated representations. In *European Conhference on Computer Vision (ECCV)*, pages 597–614. Springer, 2022. 3
- [36] David Novotny, Ignacio Rocco, Samarth Sinha, Alexandre Carlier, Gael Kerchenbaum, Roman Shapovalov, Nikita Smetanin, Natalia Neverova, Benjamin Graham, and Andrea Vedaldi. Keytr: Keypoint transporter for 3d reconstruction of deformable objects in videos. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3, 4
- [37] Pablo Palafox, Aljaz Bozic, Justus Thies, Matthias Nießner, and Angela Dai. Neural parametric models for 3d deformable shapes. In *Int. Conference on Computer Vision* (*ICCV*), 2021. 3
- [38] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. Npms: Neural parametric models for 3d deformable shapes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12695–12705, 2021. 3
- [39] Pablo Palafox, Nikolaos Sarafianos, Tony Tung, and Angela Dai. Spams: Structured implicit parametric models. In Conference on Computer Vision and Pattern Recognition (CVPR), pages 12851–12860, 2022. 3

- [40] 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 *Int. Conference on Computer Vision (ICCV)*, pages 5845– 5854. IEEE, 2021. 1, 2, 5, 6
- [41] 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 Transactions on Graphics, 40(6), 2021. 1, 2, 5, 6
- [42] Sungheon Park, Minjung Son, Seokhwan Jang, Young Chun Ahn, Ji-Yeon Kim, and Nahyup Kang. Temporal interpolation is all you need for dynamic neural radiance fields. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 4212–4221, 2023. 2
- [43] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Conference* on Computer Vision and Pattern Recognition (CVPR), pages 9054–9063, 2021. 3
- [44] Sergey Prokudin, Qianli Ma, Maxime Raafat, Julien Valentin, and Siyu Tang. Dynamic point fields. In *Int. Conference on Computer Vision (ICCV)*, 2023. 3
- [45] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10318–10327. Computer Vision Foundation / IEEE, 2021. 1, 2, 5, 6
- [46] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In Conference on Computer Vision and Pattern Recognition (CVPR), 2016. 3
- [47] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. 3
- [48] Ruizhi Shao, Zerong Zheng, Hanzhang Tu, Boning Liu, Hongwen Zhang, and Yebin Liu. Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16632–16642, 2023. 2
- [49] Samarth Sinha, Roman Shapovalov, Jeremy Reizenstein, Ignacio Rocco, Natalia Neverova, Andrea Vedaldi, and David Novotny. Common pets in 3d: Dynamic new-view synthesis of real-life deformable categories. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4881– 4891, 2023. 2, 3
- [50] Liangchen Song, Xuan Gong, Benjamin Planche, Meng Zheng, David Doermann, Junsong Yuan, Terrence Chen, and Ziyan Wu. Pref: Predictability regularized neural motion fields. In *European Configure on Computer Vision* (ECCV), pages 664–681. Springer, 2022. 2
- [51] 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. pages 2732–2742. IEEE, 2023. 2

- [52] Shih-Yang Su, Timur Bagautdinov, and Helge Rhodin. Npc: Neural point characters from video. In *Int. Conference on Computer Vision (ICCV)*, 2023. 3
- [53] Jeff Tan, Gengshan Yang, and Deva Ramanan. Distilling neural fields for real-time articulated shape reconstruction. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 4692–4701, 2023. 2, 3
- [54] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision–ECCV* 2020: 16th European Conference, Glasgow, UK, August 23– 28, 2020, Proceedings, Part II 16, pages 402–419. Springer, 2020. 4
- [55] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Nonrigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Int. Conference on Computer Vision (ICCV)*, pages 12939– 12950. IEEE, 2021. 2
- [56] Edith Tretschk, Vladislav Golyanik, Michael Zollhöfer, Aljaz Bozic, Christoph Lassner, and Christian Theobalt. Scenerflow: Time-consistent reconstruction of general dynamic scenes. In *Int. Conference on 3D Vision*, 2024. 2
- [57] Lukas Uzolas, Elmar Eisemann, and Petr Kellnhofer. Template-free articulated neural point clouds for reposable view synthesis. *arXiv pre-print*, 2023. 3
- [58] Chaoyang Wang, Lachlan Ewen MacDonald, Laszlo A Jeni, and Simon Lucey. Flow supervision for deformable nerf. In Conference on Computer Vision and Pattern Recognition (CVPR), pages 21128–21137, 2023. 2
- [59] Feng Wang, Sinan Tan, Xinghang Li, Zeyue Tian, Yafei Song, and Huaping Liu. Mixed neural voxels for fast multiview video synthesis. In *Int. Conference on Computer Vision* (*ICCV*), pages 19706–19716, 2023. 2
- [60] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Fourier plenoctrees for dynamic radiance field rendering in real-time. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13524–13534, 2022.
- [61] Liao Wang, Qiang Hu, Qihan He, Ziyu Wang, Jingyi Yu, Tinne Tuytelaars, Lan Xu, and Minye Wu. Neural residual radiance fields for streamably free-viewpoint videos. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), pages 76–87, 2023. 2
- [62] 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. In *Int. Conference on Neural Information Processing Systems (NeurIPS)*, pages 27171–27183, 2021. 2
- [63] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *Int. Conference on Computer Vision (ICCV)*, pages 3295–3306, 2023. 2
- [64] 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. 5

- [65] Christopher Wewer, Eddy Ilg, Bernt Schiele, and Jan Eric Lenssen. Simnp: Learning self-similarity priors between neural points. In Int. Conference on Computer Vision (ICCV), 2023. 3
- [66] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Wang Xinggang. 4d gaussian splatting for real-time dynamic scene rendering. arXiv pre-print, 2023. 3, 5
- [67] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9421–9431, 2021. 2
- [68] 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. 3
- [69] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Pointbased neural radiance fields. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [70] Zhen Xu, Sida Peng, Haotong Lin, Guangzhao He, Jiaming Sun, Yujun Shen, Hujun Bao, and Xiaowei Zhou. 4k4d: Real-time 4d view synthesis at 4k resolution. 2023. 2
- [71] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *Conference on Computer Vision and Pattern Recognition* (CVPR), 2022. 3
- [72] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for highfidelity monocular dynamic scene reconstruction. arXiv preprint, 2023. 3, 5
- [73] Zhengming Yu, Wei Cheng, Xian Liu, Wayne Wu, and Kwan-Yee Lin. Monohuman: Animatable human neural field from monocular video. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16943–16953, 2023. 3
- [74] Raza Yunus, Jan Eric Lenssen, Michael Niemeyer, Yiyi Liao, Christian Rupprecht, Christian Theobalt, Gerard Pons-Moll, Jia-Bin Huang, Vladislav Golyanik, and Eddy Ilg. Recent trends in 3d reconstruction of general non-rigid scenes. In *Computer Graphics Forum*. Blackwell-Wiley, 2024. 2
- [75] Qiang Zhang, Seung-Hwan Baek, Szymon Rusinkiewicz, and Felix Heide. Differentiable point-based radiance fields for efficient view synthesis. In ACM SIGGRAPH, pages 1– 12, 2022. 3
- [76] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 5
- [77] 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, 2022. 3