

3D Geometry-aware Deformable Gaussian Splatting for Dynamic View Synthesis

Zhicheng Lu^{1*}, Xiang Guo^{1*}, Le Hui^{1†}, Tianrui Chen^{1,2},
Min Yang², Xiao Tang², Feng Zhu², Yuchao Dai^{1†}

¹Northwestern Polytechnical University ²Samsung R&D Institute

{zhichenglu, guoxiang, cherryxchen}@mail.nwpu.edu.cn

{daiyuchao, huile}@nwpu.edu.cn {min16.yang, xiao1.tang, f15.zhu}@samsung.com

Abstract

In this paper, we propose a 3D geometry-aware deformable Gaussian Splatting method for dynamic view synthesis. Existing neural radiance fields (NeRF) based solutions learn the deformation in an implicit manner, which cannot incorporate 3D scene geometry. Therefore, the learned deformation is not necessarily geometrically coherent, which results in unsatisfactory dynamic view synthesis and 3D dynamic reconstruction. Recently, 3D Gaussian Splatting provides a new representation of the 3D scene, building upon which the 3D geometry could be exploited in learning the complex 3D deformation. Specifically, the scenes are represented as a collection of 3D Gaussian, where each 3D Gaussian is optimized to move and rotate over time to model the deformation. To enforce the 3D scene geometry constraint during deformation, we explicitly extract 3D geometry features and integrate them in learning the 3D deformation. In this way, our solution achieves 3D geometry-aware deformation modeling, which enables improved dynamic view synthesis and 3D dynamic reconstruction. Extensive experimental results on both synthetic and real datasets prove the superiority of our solution, which achieves new state-of-the-art performance. The project is available at <https://npucvr.github.io/GaGS/>.

1. Introduction

Dynamic View Synthesis (DVS) aims at rendering novel photorealistic views at arbitrary viewpoints and any input time step given a monocular video of a *dynamic* scene, which has broad applications in virtual reality and augmented reality. Recently, empowered with effective representations such as neural radiance fields (NeRF) [30] and Gaussian Splatting [21], novel view synthesis for static scenes has been greatly advanced. However, this success

*Equal contributions.

†Corresponding authors.

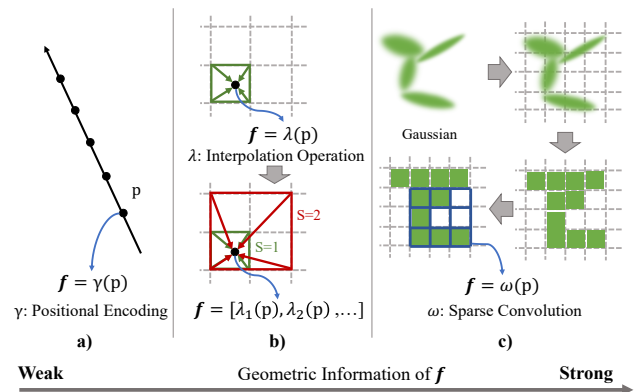


Figure 1. **Geometric information exploited by different methods.** a) Early dynamic NeRF methods such as DNeRF[37] directly encode the coordinate \mathbf{p} of the sample point as input feature for deformation network. b) Interpolation is used to fuse features from neighbouring grids and multiscale interpolation enhances the local geometry information [11, 16, 27, 53]. c) We propose to voxelize a set of Gaussian distributions and use a sparse convolution network to extract geometry-aware features for deformation learning.

cannot be extended to its dynamic counterpart directly. This is mainly due to the difficulty in modeling and representing the scene deformation. Due to the inherent motion/shape ambiguity in monocular dynamic 3D representation, dynamic scene modeling and synthesis are more challenging, especially for monocular video with limited observations.

In addressing the above challenges, one common strategy is to represent the dynamic scenes as a combination of a static canonical field and a deformation model [11, 16, 17, 27, 33, 34, 37, 51, 53], whereas the bottleneck lies in representing the diverse and complex real-world 3D deformation. To represent geometrically consistent 3D deformation, the local geometric/structural information is critical, since the deformations of the objects in the real world are highly correlated to their 3D structures. Furthermore, the motions of the object points are deeply coupled with the motions of their neighboring points. Thus, how to incorporate the local geometric information to learn locally smooth and consistent 3D deformations becomes the research focus in DVS.

Recently, different deformation models have utilized the local geometric information, but they all have their limitations. As shown in Fig. 1 a), originally in D-NeRF [37], the feature (positional encoding) of each sampled point is extracted *independently* with each other. Following works notice that this method could not handle the complex dynamic scene since the extracted features contain little information from neighboring points. In Fig. 1 b), interpolation is introduced to fuse features of neighboring grids. NDVG [16] and RoDynRF [27] gradually increase the voxel resolution so that the large voxel size could cover a larger area, introducing the local smoothness at the early stage of the training. However, this strategy has a limited cover range of local areas and cannot work at a later training stage. TiNeuVox [11] and SUDS [53] interpolate with multi-scales. Nevertheless, the interpolation operation is rather simple in extracting local geometric information and introduces un-smoothness and artifacts [3, 19].

In modling the nonrigid deformation, it is crucial to account for the consistency in the motion of local neighborhood. Since point-level MLP has a limited receptive field, which cannot capture the local geometric features of point clouds. To utilize the local geometric information effectively, we propose to use 3D sparse convolution. As shown in Fig. 1 c), building upon the recent explicit point cloud based Gaussian Splatting representation, we introduce a sparse convolution network to extract 3D geometry-aware features. Compared with simple feature interpolation, the convolutional neural network is superior in extracting local information and has a much larger reception field. Also, we treat the 3D Gaussian distributions as point clouds, which enable sparse 3D convolution for time and memory efficiency. Note that FDNerf [17] uses a 3D U-Net to inpaint the missing area in the voxel grid. But this inpaint network is not used for deformation modeling, while the rendering speed and voxel resolution are also limited.

Originally in Gaussian Splatting [21], the rotation parameter of each Gaussian is represented by quaternion. However, quaternion representation for rotation is discontinuous in parameter space for neural network learning [68]. We introduce the continuous 6D rotation [68] to ensure that the network learns a continuous function in the parameter space, which accurately represents the rotational states of each Gaussian at different time.

Overall, our method mainly has two components: a Gaussian canonical field and a deformation field. The Gaussian canonical field consists of 3D Gaussian distributions and a geometry-aware feature learning network. The explicit 3D Gaussian distribution represents the geometry of the canonical scene, and the sparse 3D CNN network extracts local structural/geometric information for each Gaussian. The deformation field estimates a transformation for each Gaussian in the canonical field, which transfers the

Gaussian from the canonical field to the given timestamp. Finally, we use 3D Gaussian splatting to render images for the given timestamp.

Our main contributions are summarized as:

- We propose a geometry-aware feature extraction network based on 3D Gaussian distribution to better utilize local geometric information.
- We propose to use continuous 6D rotation representation and modified density control strategy to adapt Gaussian splatting to dynamic scenes.
- Extensive experiments on both synthetic and real datasets show that our method surpasses competing methods by a wide margin.

2. Related Work

2.1. Novel View Synthesis

Novel View Synthesis (NVS) is a well-known task in both computer vision and graphics [5, 8, 15, 22]. Surveys such as [43, 48, 49] provide comprehensive discussions. Explicit NVS methods generally reconstruct an explicit 3D model of a scene in the form of point clouds [1], voxels, or meshes [18, 41, 42, 50]. Once the geometry of the scene is represented, novel view images can be rendered from arbitrary viewpoints via manipulating the camera pose parameters. Other methods [9, 12, 20, 35, 41, 42, 59] tackle NVS by estimating depth maps using multi-view geometry, whereas the features are aggregated from co-visible frames.

Neural Radiance Fields (NeRF) [30] is a groundbreaking approach that utilizes Multi-Layer Perceptrons (MLPs) to represent scenes implicitly. This methodology enables the modeling of a 5D radiance field, resulting in the impressive synthesis of views for static scenes. Numerous subsequent works expand the capabilities of NeRF by adapting it to various scenarios, such as handling larger and unbounded scenes [29, 40, 47, 58, 65], scene editing and relighting, [4, 45, 60, 67], [2, 3, 19], and improving the generalization ability [7, 52, 55, 64]. Meanwhile, researchers focus on achieving more efficient rendering and optimization in a NeRF-like framework. [6, 13, 24, 26, 28, 32, 36, 63] investigate efficient sampling methods along each ray for color accumulation, while [38, 39] partition the scene into multiple sub-regions as an efficient pre-processing, and [6, 13, 31, 46, 62] exploit voxel-grid representation to speed up the optimization. Very recently, [21] proposes to use 3D Gaussian distribution to represent the scene, obtaining promising results. However, these methods are mainly applicable to static scenes, and fail in scenes with dynamic objects.

2.2. Dynamic View Synthesis

A recent trend in NVS is to extend the success in static NVS to dynamic NVS. One viable strategy is to construct a 4D spatial-temporal representation. Yoon *et al.* [61] combine

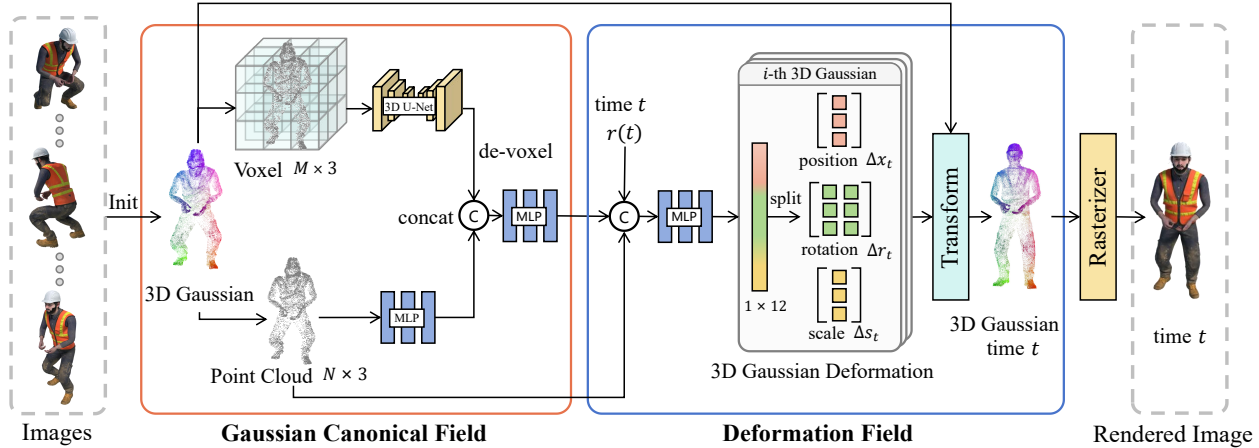


Figure 2. The pipeline of our proposed 3D geometry-aware deformable Gaussian splitting. In the Gaussian canonical field, we reconstruct a static scene in canonical space using 3D Gaussian distributions. We extract positional features using an MLP, as well as local geometric features using a 3D U-Net, fused by another MLP to form the geometry-aware features. In the deformation field, taking the geometry-aware features and timestamp t , an MLP estimates the 3D Gaussian deformation, which transfers the canonical 3D Gaussian distributions to timestamp t . Finally, a rasterizer renders the transformed 3D Gaussian to images.

single-view and multi-view depth to achieve NVS by 3D warping. Gao *et al.* [14] use a time-invariant model and a time-varying model to represent the static part and dynamic part of a scene, respectively, and use scene flow for motion modeling. NeRFlow [10] proposes a 4D spatial-temporal representation of a dynamic scene. Xian *et al.* [57] map a spatial-temporal location to the color and volume density by a 4D spatial-temporal radiance field. NSFF [23] represents a dynamic scene as a continuously changing function, encompassing various aspects of the scene, including appearance, geometry, and 3D scene motion. DCT-NeRF [54] uses the Discrete Cosine Transform (DCT) to replace the scene flow in NSFF [23] to enable smoother motion trajectories. HexPlane [6] and K-Plane [13] project 4D spatial-temporal space to multiple 2D planes.

On the other hand, works such as [11, 16, 17, 27, 33, 34, 37, 44, 51, 53] decode the dynamic scene with a canonical field and a deformation field. Along this pipeline, D-NeRF [37] first proposes the canonical-based framework. However, the deformation network utilizes positional features with little geometry information, which cannot handle complex dynamic scenarios well. Nerfies [33] proposes a coarse-to-fine optimization method for coordinate-based models that allows for more robust optimization. HyperNeRF [34] lifts the canonical field into a higher dimensional space to handle topological changes. NDVG [16] and RoDynRF [27] gradually increase the voxel resolution, which has two benefits. TiNeuVox [11] and SUDS [53] interpolate the features with multi-scales for deformation learning. The multi-scales interpolation covers a larger reception field, which benefits modeling varying motions.

Very recently, with the stunning debut of 3D Gaussian [21], some works introduce this point-based representation into their pipelines to synthesize high-fidelity images

of a dynamic scene. Wu *et al.* [56] introduce a 4D Gaussian Splatting representation and utilize a deformation field to model both Gaussian motions and shape changes. However, the multi-scale HexPlane interpolation has limited ability in extracting the geometry information, which is still insufficient for modeling complex motions. The projection-based representation compresses the 3D space to 2D space, losing 3D geometric information for deformation learning. In contrast, our canonical-based method can fully leverage the 3D information in 3D space.

3. Method

In this section, we present our 3D geometry-aware deformable Gaussian Splatting solution for dynamic view synthesis, where an overview of our pipeline is illustrated in Fig. 2. Given a set of images or monocular video of a dynamic scene with frames with corresponding time labels and known camera intrinsic and extrinsic parameters, our goal is to synthesize a novel view at any desired view at any desired time. Our method mainly consists of two core components: the Gaussian canonical field is used to learn the reconstruction of static scenes, while the deformation field is used to learn object deformation. First, we review the static 3D Gaussian splatting in Sec. 3.1. Then, we introduce the proposed Gaussian canonical field in Sec. 3.2, which consists of 3D Gaussian distributions and a geometry feature learning network. Next, in Sec. 3.3, we propose a 3D geometry-aware deformation field to learn transformations for given time steps, which transform our canonical 3D Gaussian distributions to corresponding times. In Sec. 3.4, we explain the process of rendering images from transformed 3D Gaussian distributions. Finally, we present our losses and density control modifications in Sec. 3.5.

3.1. Preliminary

3D-GS [21] represents the scene with sparse 3D Gaussians distributions. Each Gaussian has an anisotropic covariance $\Sigma \in \mathbb{R}^{3 \times 3}$ and a mean value $\mu \in \mathbb{R}^3$:

$$\mathbf{G}(\mathbf{x}) = e^{-\frac{1}{2}(\mathbf{x}-\mu)^\top \Sigma^{-1}(\mathbf{x}-\mu)}. \quad (1)$$

The covariance matrix Σ can be decomposed into a scaling matrix $\mathbf{S} \in \mathbb{R}^{3 \times 3}$ and a rotation matrix $\mathbf{R} \in \text{SO}(3)$. This ensures that the covariance matrix is positive semi-definite, while reducing the learning difficulty of 3D Gaussians:

$$\Sigma = \mathbf{R} \mathbf{S} \mathbf{S}^\top \mathbf{R}^\top. \quad (2)$$

To render an image from a designated viewpoint, the covariance matrix Σ' in camera coordinates can be calculated by giving a viewing transformation \mathbf{W} , followed by [69]:

$$\Sigma' = \mathbf{J} \mathbf{W} \Sigma \mathbf{W}^\top \mathbf{J}^\top, \quad (3)$$

where \mathbf{J} is the Jacobian of the affine approximation of the projective transformation, and \mathbf{W} is the world to camera transformation matrix.

Each Gaussian is parameterized into the following attributes: position $\mathbf{x} \in \mathbb{R}^3$, color defined by spherical harmonics coefficients $\mathbf{c} \in \mathbb{R}^k$, rotation $\mathbf{r} \in \mathbb{R}^4$, scale $\mathbf{s} \in \mathbb{R}^3$, and opacity $o \in \mathbb{R}$. Point-based α -blending and volumetric rendering like NeRF [30] essentially share the same image formation model for the splatting process. Specifically, the color \mathbf{C} of each pixel is influenced by the related Gaussians:

$$\mathbf{C} = \sum_{i=1}^N \mathbf{T}_i \alpha_i \mathbf{c}_i, \quad (4)$$

where α_i represents the density of the Gaussian point computed by a Gaussian with covariance Σ multiplied by its opacity.

3.2. Gaussian Canonical Field

In this section, we first reconstruct a static scene in canonical space. Then, we propose a geometric branch, which enables geometry feature learning of the 3D Gaussian distributions for the subsequent deformation field.

Gaussian parameters. Similar to 3D-GS [21], each Gaussian in the canonical space is characterized by position, color, scale, and opacity. Note that for rotation, we are inspired by [68] to use a continuous 6D rotation representation. Compared with the quaternion representation used in 3D-GS, the 6D rotation representation can benefit our method in estimating the deformation of each Gaussian from canonical space to time-space, especially in helping the neural networks to learn smooth rotation variation from time to time. Specifically, we set learnable parameter $[a_1, a_2]$ for each Gaussian to denote its rotation in canonical space, where a_1 and a_2 are the column vectors of three

rows, respectively. They are initialized to $[1 \ 0 \ 0]^\top$ and $[0 \ 1 \ 0]^\top$, corresponding precisely to the identity rotation matrix. The mapping from this 6D representation vector to $\text{SO}(3)$ matrix is defined as [68]:

$$f_{\text{V2M}} \left(\begin{bmatrix} | & | \\ a_1 & a_2 \\ | & | \end{bmatrix} \right) = \begin{bmatrix} | & | & | \\ b_1 & b_2 & b_3 \\ | & | & | \end{bmatrix}, \quad (5)$$

$$b_i = \left[\begin{cases} \mathcal{N}(a_1) & \text{if } i = 1 \\ \mathcal{N}(a_2 - (b_1 \cdot a_2)b_1) & \text{if } i = 2 \\ b_1 \times b_2 & \text{if } i = 3 \end{cases} \right]^\top, \quad (6)$$

where $\mathcal{N}(\cdot)$ denotes a normalization function. “ \cdot ” represents the inner product of a vector and “ \times ” represents vector cross product. V2M in f_{V2M} means the transform from 6D vector to rotation matrix.

Geometry feature learning. To capture the local geometric structure of the canonical scene, we regard the 3D Gaussian as the 3D point cloud, *i.e.*, we only use the 3D coordinates of the 3D Gaussian. In order to handle a large number of point clouds, we leverage a simple two-branch structure: the geometric branch learns local features of point clouds across different receptive fields, while the identity branch preserves the independent point-level features at high resolution. By integrating the geometric branch and identity branch, we can efficiently obtain point-level features at high resolution while embedding the local geometric information of the point cloud.

The geometric branch leverages the sparse convolution [25] on the sparse voxels to extract local geometric features at different receptive fields. Given the point cloud $\mathbf{P} \in \mathbb{R}^{N \times 3}$, we first transform the high-resolution point clouds into low-resolution voxels by dividing the space through fixed grid size s :

$$\mathbf{V} = \text{floor}(\mathbf{P}/s), \quad (7)$$

where the size of \mathbf{V} is $M \times 3$ and M is the number of voxels. Then, we construct a sparse 3D U-Net by stacking a set of sparse convolutions with a skip connection. Taking \mathbf{V} as input, we perform sparse 3D U-Net to aggregate local features (dubbed as $\mathbf{F}_v \in \mathbb{R}^{M \times C}$) of the point clouds.

The identity branch uses a multi-layer perception (MLP) to map the 3D coordinate of the point cloud into the embedding space (dubbed as $\mathbf{F}_p \in \mathbb{R}^{N \times C}$) to maintain the independence of point features. To accurately characterize the local geometric structure of the canonical scene, we fuse the voxel features with local information onto point features. Specifically, we transform the voxel feature \mathbf{F}_v back to the corresponding points to obtain point-level features $\mathbf{F}'_p \in \mathbb{R}^{N \times C}$ by assigning the voxel features to the corresponding points within it. Finally, we concatenate \mathbf{F}'_p

and \mathbf{F}_p to obtain the fused point-level feature followed by an MLP layer as:

$$\mathbf{F}_{\text{fuse}} = \text{MLP}(\text{Concat}(\mathbf{F}'_p, \mathbf{F}_p)). \quad (8)$$

3.3. Deformation Field

In this section, we propose a deformation field that estimates the deformation of each 3D Gaussian in the canonical space based on a given time t .

Deformation estimation. We adopt an MLP as the decoder \mathcal{G}_Φ , which takes the geometry feature learned from the geometry branch in the Gaussian canonical field, the position of each Gaussian, and timestamp as input, outputs the deformation of each Gaussian from canonical space to time t , including position deformation $\Delta\mathbf{x}_t \in \mathbb{R}^3$, rotation deformation $\Delta\mathbf{r}_t \in \mathbb{R}^6$ and scale deformation $\Delta\mathbf{s}_t \in \mathbb{R}^3$:

$$\Delta\mathbf{x}_t, \Delta\mathbf{r}_t, \Delta\mathbf{s}_t = \mathcal{G}_\Phi(\mathbf{F}_{\text{fuse}}, \gamma(\mathbf{x}), \gamma(t)), \quad (9)$$

where $\gamma(\cdot)$ denotes the positional encoding in NeRF [30], which maps a one dimension signal from \mathbb{R} into a higher dimensional space \mathbb{R}^{2L} :

$$\begin{aligned} \gamma(p) = & (\sin(2^0\pi p), \cos(2^0\pi p), \\ & \dots, \\ & \sin(2^{L-1}\pi p), \cos(2^{L-1}\pi p)). \end{aligned} \quad (10)$$

Note that we set the color parameters \mathbf{c} and opacity o of canonical 3D Gaussian distributions constant over time. These two factors are highly related to the physical properties of the Gaussian distributions, and we want each distribution to represent the same object area over the timeline.

Transformation. Using the estimated deformation for time t above, we could transform the 3D Gaussian distributions to current time by

$$\begin{aligned} \mathbf{x}_t &= \mathbf{x} + \Delta\mathbf{x}_t, \\ \mathbf{s}_t &= \mathbf{s} + \Delta\mathbf{s}_t, \\ \mathbf{r}_t &= f_{\text{V2M}}(\Delta\mathbf{r}_t) \times f_{\text{V2M}}(\mathbf{r}). \end{aligned} \quad (11)$$

3.4. Rasterization

Once we have completed preparing the attributes of each Gaussian $(\mathbf{x}_t, \mathbf{c}, \mathbf{r}_t, \mathbf{s}_t, o)$, we use the differentiable tile rasterizer [21] to render the image at any desired viewpoint at this timestamp:

$$\hat{\mathbf{C}}_t = \text{Rasterizer}(\mathbf{x}_t, \mathbf{c}, \mathbf{r}_t, \mathbf{s}_t, o, \mathbf{K}, [\mathbf{R}|\mathbf{T}]), \quad (12)$$

where \mathbf{K} and $[\mathbf{R}|\mathbf{T}]$ represent the camera's intrinsic and extrinsic parameters, respectively.




Figure 3. Our density control is designed for dynamic scenes. We control the densification of Gaussian distributions according to their transformed parameters at timestamp t rather than parameters at canonical space.

3.5. Optimization

To optimize the model, we use the photometric loss, and a motion loss, and also adapt the density control from 3D-GS [21] with our modifications.

Photometric loss. The photometric loss consists of the L_1 loss and structural similarity loss L_{D-SSIM} between the rendered image $\hat{\mathbf{C}}_t$ and ground truth image \mathbf{C}_t .

$$L_{\text{photo}} = (1 - \lambda)L_1 + \lambda L_{D-SSIM}. \quad (13)$$

Regularization. We accept the fact that in a scene, the proportion of dynamic points is much smaller than that of static points, and the motion amplitude at dynamic points is not too large. In other words, the point in a scene should be as static as possible,

$$L_{\text{motion}} = \|\Delta\mathbf{x}_t\|_1. \quad (14)$$

Total loss. The total loss we used is defined as follows,

$$L = L_{\text{photo}} + \omega L_{\text{motion}}, \quad (15)$$

where ω is a trade-off parameter to balance the components.

Density control. 3D-GS has shown that adaptive density control is essential in achieving high rendering performance. On the one hand, the Gaussians need to populate empty areas without geometric features. Thus, it simply creates a copy of the Gaussian for under-reconstructed regions. On the other hand, large Gaussians in regions with high variance need to be split into smaller Gaussians. We implement our method like 3D-GS but replace such Gaussians with two new ones, divide their scale by a factor of $\phi = 1.6$, and initialize their position by using the original 3D Gaussian as a PDF for sampling.

Our method differs from 3D-GS in the following aspects. For 3D-GS, there only exists sets of Gaussians. However, in our case, we initialize the Gaussians in the canonical space, then estimate the deformations of these Gaussians,

Table 1. Quantitative comparison between our method and competing methods on the D-NeRF dataset. The best results are highlighted in bold.

Method	Hell Warrior			Mutant			Hook			Bouncing Balls		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
3D-GS [21]	15.3924	0.8776	0.1300	21.7554	0.9359	0.0575	18.6933	0.8733	0.1144	22.5575	0.9485	0.0647
D-NeRF [37]	25.0293	0.9506	0.0691	31.2900	0.9739	0.0268	29.2567	0.9650	0.1174	38.9300	0.9900	0.1031
TiNeuVox-B[11]	28.2058	0.9661	0.0631	33.9029	0.9771	0.0301	31.7929	0.9718	0.0436	40.8536	0.9913	0.0401
NDVG [16]	26.4933	0.9600	0.0670	34.4131	0.9801	0.0270	30.0009	0.9626	0.0463	37.5157	0.9874	0.0751
FDNeRF [17]	27.7120	0.9665	0.0508	34.9727	0.9810	0.0312	32.2867	0.9756	0.0388	40.0191	0.9912	0.0395
4D-GS [56]	28.1196	0.9730	0.0276	38.3411	0.9936	0.0062	33.1560	0.9810	0.0168	40.7418	0.9941	0.0105
Ours	32.2712	0.9835	0.0164	41.4284	0.9969	0.0029	36.9647	0.9916	0.0076	43.5929	0.9960	0.0061

Method	Lego			T-Rex			Stand Up			Jumping Jacks		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
3D-GS [21]	23.0991	0.9329	0.0567	25.7496	0.9567	0.0474	19.3779	0.9200	0.0909	20.7163	0.9227	0.0980
D-NeRF [37]	21.6427	0.8394	0.1654	31.7568	0.9767	0.0396	32.7992	0.9818	0.0215	32.8031	0.9810	0.0373
TiNeuVox-B[11]	25.1748	0.9217	0.0689	32.7750	0.9783	0.0307	36.2031	0.9859	0.0199	34.7390	0.9823	0.0328
NDVG [16]	25.0416	0.9395	0.0534	32.6229	0.9781	0.0330	33.2158	0.9793	0.0302	31.2530	0.9737	0.0398
FDNeRF [17]	25.2700	0.9390	0.0460	30.7068	0.9731	0.0368	36.9107	0.9878	0.0188	33.5521	0.9812	0.0329
4D-GS [56]	25.4024	0.9434	0.0377	33.3912	0.9869	0.0130	38.2610	0.9923	0.0071	35.6656	0.9882	0.0159
Ours	25.4411	0.9474	0.0329	39.0285	0.9952	0.0052	42.2101	0.9966	0.0028	37.9604	0.9928	0.0088

Table 2. Quantitative comparison between our method and competing methods on the HyperNeRF dataset. The best results are highlighted in bold.

Method	Chicken		3D Printer		Broom		Peel Banana	
	PSNR \uparrow	MS-SSIM \uparrow	PSNR \uparrow	MS-SSIM \uparrow	PSNR \uparrow	MS-SSIM \uparrow	PSNR \uparrow	MS-SSIM \uparrow
TiNeuVox[11]	28.2861	0.9474	22.7514	0.8392	21.2682	0.6832	24.5136	0.8743
NDVG [16]	27.0536	0.9390	22.4196	0.8389	21.4658	0.7028	22.8204	0.8279
FDNeRF [17]	27.9627	0.9438	22.8027	0.8453	21.9091	0.7154	24.2515	0.8645
3D-GS [21]	20.8915	0.7426	18.3991	0.6114	20.3953	0.6598	20.5654	0.8094
Ours	28.5342	0.9331	22.0403	0.8098	20.8994	0.5241	25.5785	0.9067

Table 3. Quantitative comparison on HyperNeRF dataset: Average on Cut Lemon, Chicken, 3D Printer, and Split Cookie. The best results are highlighted in bold.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
TiNeuVox-B [11]	27.16	0.76	0.40
3D-GS [21]	21.26	0.69	0.40
4D-GS [56]	26.98	0.78	0.31
Ours	27.52	0.80	0.25

and transform their attributes into a timestamp space. As shown in Fig. 3, we use the Gaussians at the current moment to render the image. Therefore, we determine whether the Gaussians need to conduct density control by the current attributes (like scale) at the current timestamp rather than the canonical attributes. Afterward, we inverse the transformation of the split/cloned Gaussian back to the canonical space.

4. Experiments

4.1. Dataset

In the paper, we use both synthetic and real datasets for evaluating our method. The synthetic dataset D-NeRF [37] contains 8 dynamic scenes, including Hell Warrior, Mutant, Hook, Bouncing Balls, Lego, T-Rex, Stand Up, and Jumping Jacks. The real dataset proposed by HyperNeRF [34], including interp-cut-lemon, interp-cut-lemon1, vrig-

chicken, vrig-3dprinter, misc-split-cookie, and misc-split-cookie. Following previous works [21], we report three evaluation metrics, including Peak Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [66].

4.2. Implementation Details

Our implementation is based on 3D-GS [21]. We trained a total of 40000 iterations, with the first 3000 iterations only optimizing static scenes, and then adding deformation fields to optimize dynamic scenes. The learning rate of our network takes an exponential decay from $8e-4$ to $1.6e-6$ with the Adam optimizer. Moreover, we use a 2-layer MLP with a width of 64 for the front point feature extraction, and a 3-layer MLP with a width of 64 for the back point feature fusion. Then 5 layers MLP with width 256 and skip connection is used for a decoder. For the positional encoding process, we use $L = 10$ for position \mathbf{x} and $L = 6$ for timestamp t . For the D-NeRF dataset, which does not provide point clouds, we randomly initialize 150000 points. Meanwhile, for the HyperNeRF dataset, we use the point cloud provided in its dataset as the initial point cloud. All the experiments are tested on a single RTX 4090 GPU.

4.3. Quantitative Results

Synthetic scenes. We compare our method with recent state-of-the-art methods in the field, including 3D-GS, D-




Figure 4. Qualitative comparisons between baselines and our method on the synthetic dataset.




Figure 5. Qualitative comparisons between baselines and our method on the HyperNeRF real dataset[34].

NeRF, TiNeuVox, NDVG, FDNeRF, and 4D-GS on the D-NeRF Dataset. As shown in Table 1, we list the results of each scene. It can be observed that our method is significantly better than other methods in terms of all three metrics for physical canonical-based methods. On average, our method significantly improves PSNR compared with static Gaussian, 3D-GS. The computational costs are: training time around 2h (avg. on D-NeRF dataset), render FPS 12 (fixed viewpoint), model size (34MB points cloud + 14MB network). Since it inherently cannot model the deformation of the dynamic scene, 3D-GS performs poorly in dynamic view synthesis.

Real scenes. We further compare our method with some highly related works on the real scene dataset proposed by [34]. We have shown the detailed results on chicken, 3D printer, broom, and peel banana in Table 2, and the average result on cut lemon, chicken, 3d printer, split cookie in Table 3. It can be observed that our method achieves good performance compared with other state-of-the-art methods. Compared with synthetic datasets, real datasets are more challenging due to the narrow camera viewing range and pose ambiguity. The quantitative results can demonstrate the effectiveness of the proposed method in real scenes.

4.4. Visualization Results

Visual comparison. In addition to quantitative results, we also provide visualization results of different methods to demonstrate the superiority of our method. For better comparison, we show the rendered images of each synthetic scene from the same viewpoint in Fig. 4. By comparing the visualization results of different methods, it is shown that the rendered images by our method are closer to the ground truth images, indicating that our method can recover accurate and detailed images. In addition, we provide visualization results of the real scenes in Fig. 5. Compared with TiNeuVox [11], our method can recover the detailed structure of dynamic objects, like chicken and banana.

Gaussian visualization. To verify the effectiveness of our method, we show the 3D point cloud of the 3D Gaussian. Specifically, we only use the 3D coordinates of the 3D Gaussian. As shown in Fig. 7, we provide the point clouds of different methods on the synthetic dataset, including 3D-GS [21], 4D-GS [56], and ours. Note that the color of the point cloud is generated by 3D coordinates. Since 3D-GS cannot model dynamic scenes, the quality of the point cloud is poor. Comparing 4D-GS with ours, it can be observed




Figure 6. Visualization of learned geometry-aware features.




Figure 7. Visualization of learned Gaussian. Colored with position coordinates

that the point cloud of our method has a clear local geometric structure.

4.5. Ablation Study

We conduct ablation studies on the synthetic dataset (800×800) to verify the effectiveness of our proposed components. In Table 4, vanilla model is a simple MLP model without our components.

Effect of geometric-aware features. To learn the geometric information of the object in our Gaussian canonical field, we voxelize the 3D Gaussian distributions and extract geometric aware features using our 3D U-Net. To demonstrate the effectiveness of this design, we test our method with geometric branch blocks and leave others unchanged. In Table 4, ours full has a clear advantage over w/o geo. feat., and our geometry branch plays the most important role among the components studied in the ablations.

In Fig. 6, we visualize the learned geometric-aware features. We color the point clouds with the learned features, and it shows meaningful geometric information. Interestingly, we can see an obvious difference in the learned features between the moving objects (bucket of the lego and the t-rex body) and the static objects (body of the lego and the ground in t-rex). Also, our geometric-aware features reflect the local geometric structure. For example, the spines of the bones on the t-rex tail have similar features, and the smooth part of the tail bones have other patterns.

Different geometric features. We use the PointNet-like architecture and plane projection (2D CNN) to conduct ex-

Table 4. **Ablation Study.** Ablation studies in terms of average PSNR, SSIM, and LPIPS. The best results are highlighted in bold.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o geo. feat.	37.5757	0.9841	0.0173
w/o 6D rotation	37.8750	0.9851	0.0154
canonical DC	37.8026	0.9847	0.0166
vanilla	35.2307	0.9793	0.0242
PointNet feat.	36.7353	0.9826	0.0184
Plane feat.	35.9054	0.9811	0.0212
ours full	38.0134	0.9853	0.0153

periments. Compared with the results (dubbed as “PointNet feat.” and “Plane feat.”) in Table 4, it can be observed that our method achieves significant performance gains.

6D representation. To study the effect of 6D representation of the rotation parameters of the 3D Gaussian, we conduct an experiment that replaces the 6D vector with quaternion \mathbf{q} which is used in the original 3D-GS. To deform the 3D Gaussian in canonical space, our deformation field estimates a $\Delta\mathbf{q}_t$ and gets $\mathbf{q}_t = \mathbf{q} + \Delta\mathbf{q}_t$, using the quaternion add operation. In Table 4, quaternion demonstrates an obvious performance drop, which proves the effectiveness of the 6D representation.

Density control. In terms of density control, we test the setting that only uses the 3D Gaussian in canonical space without considering the transform 3D Gaussian at other timestamps. In Table 4, canonical DC shows a performance drop, as the canonical 3D Gaussian alone cannot reflect the over/under reconstruction information at all timestamps for dynamic scenes.

5. Conclusion

In this paper, we have proposed a 3D geometry aware Gaussian Splatting solution for dynamic view synthesis. We addressed the limitations of existing approaches from two perspectives: 1) we introduced 3D sparse convolution to extract local structural information effectively and efficiently for deformation learning, and 2) we represented the dynamic scenes as a collection of deforming 3D Gaussian distributions, which are optimized to deform (move, rotate, scaling) over time. Experimental results across synthetic and real datasets demonstrate the superiority of our solution in dynamic view synthesis and 3D reconstruction. We plan to further investigate explicit motion modeling by exploiting the foreground and background motion segmentation cues.

Acknowledgments

We thank the area chairs and the reviewers for their insightful and positive feedback. This work was supported in part by the National Science Fund of China (Grant Nos. 62271410, 62306238) and the Fundamental Research Funds for the Central Universities.

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Zip-NeRF: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023. 2
- [4] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. NeRD: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [5] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the Conference on Computer Graphics and Interactive Techniques*, 2001. 2
- [6] Ang Cao and Justin Johnson. HexPlane: A fast representation for dynamic scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 3
- [7] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. MVSNeRF: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [8] Shenchang Eric Chen and Lance Williams. View interpolation for image synthesis. In *Proceedings of the Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1993. 2
- [9] Inchang Choi, Orazio Gallo, Alejandro Troccoli, Min H Kim, and Jan Kautz. Extreme view synthesis. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [10] Yilun Du, Yanan 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 (ICCV)*, 2021. 3
- [11] 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 *Proceedings of the Conference on Computer Graphics and Interactive Techniques in Asia (SIGGRAPH ASIA)*, 2022. 1, 2, 3, 6, 7
- [12] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. DeepStereo: Learning to predict new views from the world’s imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [13] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-Planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 3
- [14] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 3
- [15] Ned Greene. Environment mapping and other applications of world projections. *IEEE Computer Graphics and Applications*, 1986. 2
- [16] Xiang Guo, Guanying Chen, Yuchao Dai, Xiaoqing Ye, Jiadai Sun, Xiao Tan, and Errui Ding. Neural deformable voxel grid for fast optimization of dynamic view synthesis. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, 2022. 1, 2, 3, 6
- [17] 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 *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023. 1, 2, 3, 6
- [18] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (TOG)*, 2018. 2
- [19] Wenbo Hu, Yuling Wang, Lin Ma, Bangbang Yang, Lin Gao, Xiao Liu, and Yuewen Ma. Tri-MipRF: Tri-mip representation for efficient anti-aliasing neural radiance fields. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023. 2
- [20] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi. Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 2016. 2
- [21] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (TOG)*, 2023. 1, 2, 3, 4, 5, 6, 7
- [22] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1996. 2
- [23] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [24] David B Lindell, Julien NP Martel, and Gordon Wetzstein. AutoInt: Automatic integration for fast neural volume rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [25] Baoyuan Liu, Min Wang, Hassan Foroosh, Marshall Tappen, and Marianna Pensky. Sparse convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 4
- [26] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *Proceed-*

- ings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [27] Yu-Lun Liu, Chen Gao, Andreas Meuleman, Hung-Yu Tseng, Ayush Saraf, Changil Kim, Yung-Yu Chuang, Johannes Kopf, and Jia-Bin Huang. Robust dynamic radiance fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 3
- [28] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Transactions on Graphics (TOG)*, 2021. 2
- [29] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [30] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 4, 5
- [31] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 2022. 2
- [32] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Joerg H. Mueller, Chakravarty R. Alla Chaitanya, Anton S. Kaplanyan, and Markus Steinberger. DONeRF: Towards real-time rendering of compact neural radiance fields using depth oracle networks. *Computer Graphics Forum (CGF)*, 2021. 2
- [33] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 1, 3
- [34] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. HyperNeRF: A higher-dimensional representation for topologically varying neural radiance fields. *ACM Transactions on Graphics (TOG)*, 2021. 1, 3, 6, 7
- [35] Eric Penner and Li Zhang. Soft 3D reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)*, 2017. 2
- [36] Martin Píala and Ronald Clark. TerminiNeRF: Ray termination prediction for efficient neural rendering. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2021. 2
- [37] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 3, 6
- [38] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. DeRF: Decomposed radiance fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [39] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. KiloNeRF: Speeding up neural radiance fields with thousands of tiny MLPs. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [40] Konstantinos Rematas, Andrew Liu, Pratul P. Srinivasan, Jonathan T. Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [41] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [42] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [43] Harry Shum and Sing Bing Kang. Review of image-based rendering techniques. In *Visual Communications and Image Processing (VCIP)*, 2000. 2
- [44] Liangchen Song, Anpei Chen, Zhong Li, Zhang Chen, Lele Chen, Junsong Yuan, Yi Xu, and Andreas Geiger. NeRF-Player: A streamable dynamic scene representation with decomposed neural radiance fields. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 2023. 3
- [45] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [46] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct Voxel Grid Optimization: Super-fast convergence for radiance fields reconstruction. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [47] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretzschmar. Block-NeRF: Scalable large scene neural view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [48] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. State of the art on neural rendering. In *Computer Graphics Forum (CGF)*, 2020. 2
- [49] Ayush Tewari, O Fried, J Thies, V Sitzmann, S Lombardi, Z Xu, T Simon, M Nießner, E Tretschk, L Liu, et al. Advances in neural rendering. In *Proceedings of the Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 2021. 2
- [50] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 2019. 2
- [51] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Pro-*

- ceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 1, 3
- [52] Alex Trevithick and Bo Yang. GRF: Learning a general radiance field for 3D representation and rendering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [53] Haithem Turki, Jason Y Zhang, Francesco Ferroni, and Deva Ramanan. SUDS: Scalable urban dynamic scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 3
- [54] Chaoyang Wang, Ben Eckart, Simon Lucey, and Orazio Gallo. Neural trajectory fields for dynamic novel view synthesis. *arXiv preprint arXiv:2105.05994*, 2021. 3
- [55] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. IBRNet: Learning multi-view image-based rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [56] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4D Gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. 3, 6, 7
- [57] 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 (CVPR)*, 2021. 3
- [58] Yuanbo Xiangli, Linning Xu, Xingang Pan, Nanxuan Zhao, Anyi Rao, Christian Theobalt, Bo Dai, and Dahua Lin. BungeeNeRF: Progressive neural radiance field for extreme multi-scale scene rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [59] Zexiang Xu, Sai Bi, Kalyan Sunkavalli, Sunil Hadap, Hao Su, and Ravi Ramamoorthi. Deep view synthesis from sparse photometric images. *ACM Transactions on Graphics (TOG)*, 2019. 2
- [60] Wenqi Yang, Guanying Chen, Chaofeng Chen, Zhenfang Chen, and Kwan-Yee K Wong. S³-NeRF: Neural reflectance field from shading and shadow under a single viewpoint. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2
- [61] Jae Shin Yoon, Kihwan Kim, Orazio Gallo, Hyun Soo Park, and Jan Kautz. Novel view synthesis of dynamic scenes with globally coherent depths from a monocular camera. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [62] Alex Yu, Sara Fridovich-Keil, 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 (CVPR)*, 2021. 2
- [63] 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 (ICCV)*, 2021. 2
- [64] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [65] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. NeRF++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 2
- [66] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [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 Transactions on Graphics (TOG)*, 2021. 2
- [68] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 4
- [69] M. Zwicker, H. Pfister, J. van Baar, and M. Gross. Ewa volume splatting. In *Proceedings of IEEE Visualization (VIS)*, 2001. 4