

Forward Flow for Novel View Synthesis of Dynamic Scenes

Xiang Guo¹, Jiadai Sun^{1,3}, Yuchao Dai^{†1}, Guanying Chen², Xiaoqing Ye³, Xiao Tan³, Errui Ding³, Yumeng Zhang³, Jingdong Wang³

¹Northwestern Polytechnical University, ²FNii and SSE, CUHK-Shenzhen, ³Baidu Inc.

Abstract

This paper proposes a neural radiance field (NeRF) approach for novel view synthesis of dynamic scenes using forward warping. Existing methods often adopt a static NeRF to represent the canonical space, and render dynamic images at other time steps by mapping the sampled 3D points back to the canonical space with the learned backward flow field. However, this backward flow field is non-smooth and discontinuous, which is difficult to be fitted by commonly used smooth motion models. To address this problem, we propose to estimate the forward flow field and directly warp the canonical radiance field to other time steps. Such forward flow field is smooth and continuous within the object region, which benefits the motion model learning. To achieve this goal, we represent the canonical radiance field with voxel grids to enable efficient forward warping, and propose a differentiable warping process, including an average splatting operation and an inpaint network, to resolve the many-to-one and one-to-many mapping issues. Thorough experiments show that our method outperforms existing methods in both novel view rendering and motion modeling, demonstrating the effectiveness of our forward flow motion modeling. Project page: <https://npucvr.github.io/ForwardFlowDNeRF>.

1. Introduction

Novel view synthesis (NVS) is a challenging and long-standing problem in computer vision and graphics, which has many applications in virtual reality, augmented reality, data augmentation, image editing, *etc.* Recently, differentiable neural rendering [26, 30, 59] has been introduced into this area. In particular, the neural radiance field (NeRF) [26]

[†] Corresponding author (daiyuchao@nwpu.edu.cn). The first three authors also with the Shaanxi Key Laboratory of Information Acquisition and Processing.

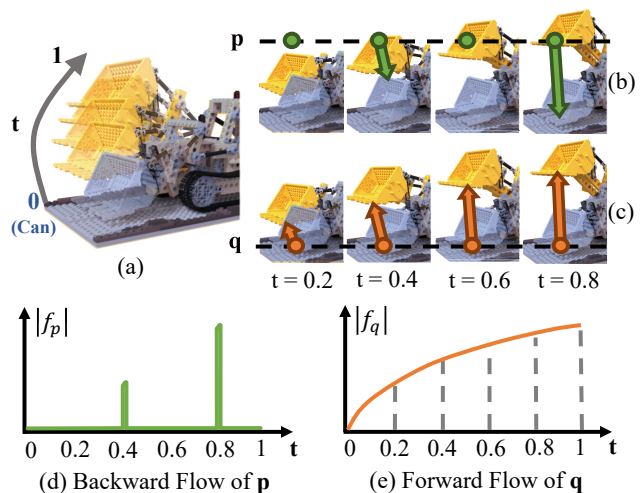


Figure 1. **Comparison of backward flow and forward flow.** This figure shows an example of backward and forward flow changes. (a) An example of dynamic scene. (b) With the bucket lifting up, different types of points cover the green point \mathbf{p} , which needs very different backward flows to map this point back to canonical space. (d) shows the norm changes of the backward flow, which is not smooth. (c) On the other hand, the forward flow of position \mathbf{q} , which maps the constant object point from canonical space to other times, is smooth and continuous. (e) shows the norm changes of the forward flow.

promotes this area significantly and attracts lots of interest within a short time. NeRF [26] produces realistic images by representing the 3D world with a multi-layer perceptron (MLP), which maps the input 3D coordinates and 2D view direction to target density and color.

While the original NeRF [26] can only model static scenes, a series of works extend the NeRF-based framework from static to dynamic scenes [9, 13, 19, 32, 36, 49, 51, 55]. One of the promising directions is using a canonical space representation [15, 36, 49]. This representation sets one of the time steps as canonical time and models the static scene with a canonical radiance field. To render images at other time steps, a deformation field is used to estimate the

backward flow for moving the 3D points from the current time step back to the canonical time step. Although the canonical-based representation with the backward flow is easy to implement, the backward flow field is non-smooth and discontinuous. As shown in Fig. 1(b), for a fixed 3D position along the timeline, there will be different types of points covering the position \mathbf{p} , which needs discontinuous flows to map them back to canonical space (Fig. 1(d)). So the backward flow cannot be fitted well with commonly used smooth motion models (MLP, for example). Also, distortions are introduced to the static canonical space geometry due to the failure of the motion model, as shown in Fig. 8.

To solve the problem of backward flow, we propose using *forward flow* as the deformation model. Instead of warping the sampled points on image rays at other time steps back to the canonical time and rendering at the canonical space, we propose to warp the whole canonical radiance field from the canonical time to other time steps using forward deformation flow and render at corresponding time steps. In this way, the forward flow estimated by the deformation model will be smooth and continuous for the same 3D position along the timeline (Fig. 1(c) and (e)). Note that SNARF [7] has also used forward warping based on an inverse skinning model, but it is designed for dynamic human modeling and cannot be used in general scenes. In this paper, we aim to achieve forward warping for general scenes, which means we must warp the whole space.

However, introducing forward warping into the canonical space based NeRF methods is not straightforward as there are three main problems to be solved. First, the traditional canonical radiance field in existing methods cannot be warped explicitly, since the radiance field is represented as a continuous function parameterized by an MLP. To solve this problem, we propose to use the voxel grid to represent the canonical radiance field as it is finite and discrete. Recent voxel-based methods [27,44,62] have proven the effectiveness of this representation. The other two problems are the *many-to-one* and *one-to-many* mapping issues brought by the inherent property of the forward warping operation. To address them, we propose a *differentiable forward warping* method consisting of an *average splatting* operation and an *inpaint network* to solve the *many-to-one* and *one-to-many* issue, respectively. Extensive experiments have been conducted to verify the effectiveness of our method.

Our key contributions can be summarized as follows:

- To the best of our knowledge, we are the first to investigate forward warping in dynamic view synthesis for general scenes. We propose a novel canonical based NeRF with forward flow motion modeling for dynamic view synthesis. Thanks to the forward flow field, our method can better represent the object motions, and explicitly recover the trajectory of a surface point.

- We introduce voxel grid based canonical radiance field to enable reasonable computation of forward warping, and propose a differentiable forward warping method, including an average splatting operation and an inpaint network, to solve the many-to-one and one-to-many issues of forward warping.
- Experiments on multiple datasets show that our method outperforms existing methods on the D-NeRF [36] dataset, Hypernerf [33] dataset, NHR [54] dataset and our proposed dataset.

2. Related Works

Novel View Synthesis (NVS) is a long-standing task in both computer vision and graphics [4, 6, 14, 18], and surveys of recent methods can be found in [42, 46, 47]. Some methods require to reconstruct an explicit 3D model to represent the scene, such as point clouds [1], voxels or meshes [16, 40, 41, 48]. Then novel images from arbitrary views can be rendered from this geometry. Another methods try to estimate depth map through multi-view geometry, and then aggregate features from multiple frames through the co-visibility, such as [8, 11, 17, 34, 40, 41, 57]. Recently, neural implicit representations have shown great promise for novel view synthesis and 3D modeling. In this section, we focus more on the neural implicit rendering methods and mainly summarize these schemes according to whether they can handle dynamic scenes.

NVS for Static Scenes NeRF [26], as a seminal work, uses MLPs to model a 5D radiance field, which can render impressive view synthesis for static scenes captured. Numerous subsequent works have extended NeRF to kinds of scenarios, such as larger and unbounded scenes [25, 39, 45, 56, 64], relighting [3, 43, 58, 66], incorporating anti-aliasing for multi-scale rendering [2], and generalization ability [5, 50, 52, 63]. In addition, some methods are devoted to more efficient neural rendering and optimization in NeRF-like framework, such as [20, 22, 24, 28, 35, 62] focus more on efficient sampling along each ray for color accumulation, [37, 38] subdivide the scene into multiple cells for efficient processing, and [27, 44, 61] exploit voxel-grid representation to speed up the optimization of radiance field. However, these methods are mainly applicable to static scenes, leaving out the scenes with dynamic objects, which are actually more extensive and practical.

NVS for Dynamic Scenes There are several works that extend NeRF from static scenes to dynamic scenes with non-rigid deformable objects. One feasible way is to build a 4D spatial-temporal representation. For example, Yoon *et al.* [60] combine single-view depth and depth from multi-view stereo to render virtual views with 3D warping. Gao *et al.* [13] use a time-invariant model (static) and a time-varying model (dynamic) to represent the scenes, and reg-

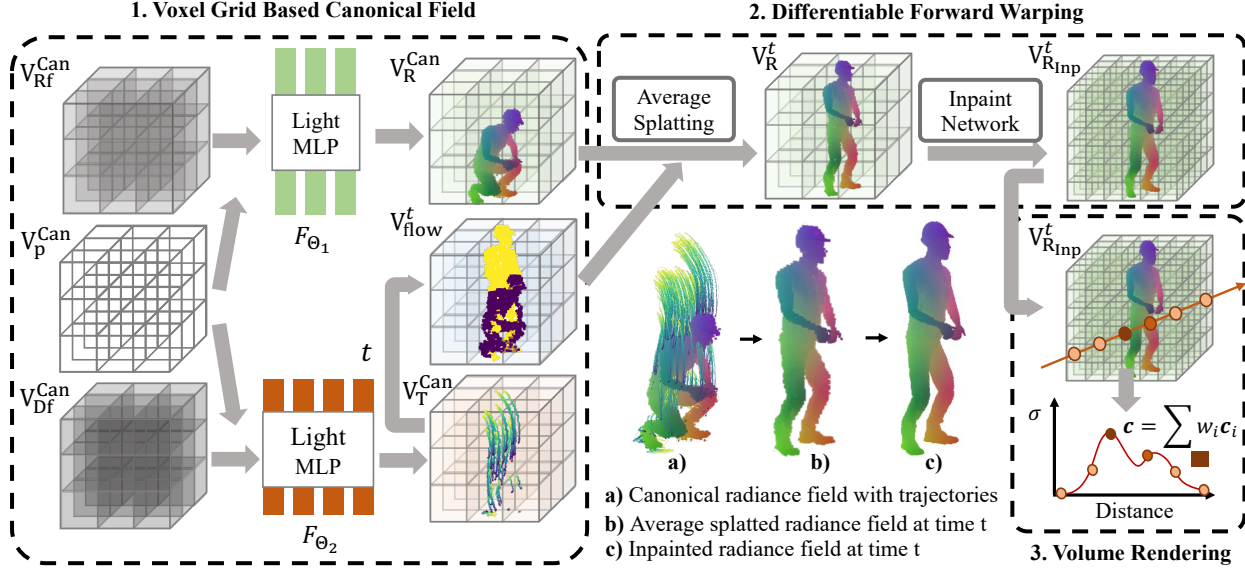


Figure 2. **Overview of our proposed method.** **a)** We represent a static scene at canonical time with a voxel grid based radiance field for density&color and a voxel grid based trajectory field for deformations; **b)** We propose to first forward warp canonical radiance field using the forward flow by average splatting; **c)** We then inpaint the warped radiance field using an inpaint network; Specifically, **1. Voxel Grid Based Canonical Field** contains two models. The canonical radiance field $\mathbf{V}_R^{\text{Can}}$ is estimated by a Light MLP which takes canonical radiance feature $\mathbf{V}_{Rf}^{\text{Can}}$ and corresponding 3D coordinates $\mathbf{V}_p^{\text{Can}}$ as input. The canonical trajectory field $\mathbf{V}_T^{\text{Can}}$ is estimated by another Light MLP which takes deformation feature and coordinates as input. The deformation flow $\mathbf{V}_{\text{flow}}^t$ from canonical to time t can then be obtained; **2. Differential Forward Warping** first warp $\mathbf{V}_R^{\text{Can}}$ to get radiance field \mathbf{V}_R^t at time t . Then, the \mathbf{V}_R^t is inpainted by an inpaint network, which is $\mathbf{V}_{R\text{Inp}}^t$; **3. Volume Rendering** render colors of rays at time t based on $\mathbf{V}_{R\text{Inp}}^t$

ularize the dynamic model by scene flow estimation. NeRF-Flow [9] learns a 4D spatial-temporal representation of a dynamic scene from a set of RGB images. Xian *et al.* [55] build a 4D space-time irradiance field to map a spatial-temporal location to the emitted color and volume density. Similarly, NSFF [19] models the dynamic scene as a time-variant continuous function of appearance, geometry, and 3D scene motion. DCT-NeRF [51] uses the Discrete Cosine Transform (DCT) to capture dynamic motion, *i.e.* learning smooth and stable trajectories over time for each point in space.

On the other hand, D-NeRF [36], Nerfies [32], HyperNeRF [33] and NR-NeRF [49] use a static canonical radiance field to capture geometry and appearance, and then learn a deformation/displacement field at each time step w.r.t. the canonical space. Specifically, to render an image at an arbitrary time step, a deformation field is used to estimate backward scene flow, moving 3D points from the current time step back to the canonical step. However, for the same 3D location along the timeline, the backward flow field is not guaranteed to be smooth and continuous. As a result, the canonical geometry usually has distortions and resembles the mean shape of a moving object. We focus on solving the problem of backward flow in this paper.

Along with the two main directions, there is a trend to speed up the training of dynamic NeRF, which is based on voxel grid representation. TiNeuVox [10] models the de-

formation using a tiny MLP and uses multi-distance interpolation to get the feature for the radiance network which estimates the density and color. V4D [12] uses the 3D feature voxel to model the 4D radiance field with additional time dimension concatenated and proposes look-up tables for pixel-level refinement. Although V4D mainly focuses on improving image quality, the training speed is not significant compared to TiNeuVox. DeVRF [21] also builds on voxel-grid representation, which proposes to use multi-view data to overcome the nontrivial problem of the monocular setup. Multi-view data simplifies the learning of motion and geometry compared with others using monocular images.

3. Method

Motivation Backward-warping based methods [15, 36] propose a network $f_{t \rightarrow \text{Can}} = F(p, t)$ to estimate the deformation flow $f_{t \rightarrow \text{Can}}$ which moves the point at the position p from other time steps t back to canonical time Can. However, for the same position p , at different time steps t s, there could be different object points or even the empty point at this position, as shown in Fig. 1. This means the deformation flow $f_{t \rightarrow \text{Can}}$ is non-smooth and discontinuous along the timeline for a fixed position p . This could introduce difficulties for motion learning and produce distortion in the canonical radiance field, because the backward flow network has limited capacities in learning a correct non-smooth and dis-

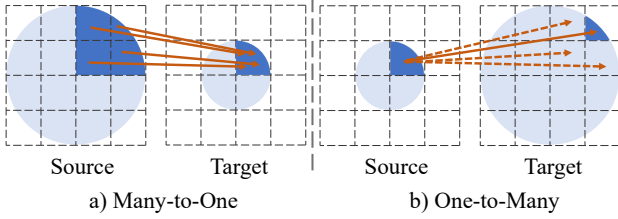


Figure 3. **Two issues of the forward warping.** Many-to-one: multiple positions in the source correspond to the same position in the target. One-to-many: one position in the source corresponds to multiple positions in the target, causing holes if warping naively.

continuous deformation field. On the other hand, our proposed forward warping strategy estimates the deformation $f_{\text{Can} \rightarrow t}$ from the canonical time to another time step of a 3D point in canonical space. The set of the deformation flows of one position along the timeline is actually the trajectory of this point. This guarantees the flows to be smooth and continuous along the timeline as we assume the motions in reality have these properties.

Method Overview We model a static scene at the canonical time using the voxel grid representation with a canonical radiance field and a canonical trajectory field. To synthesize dynamic images, we propose to forward warp the canonical radiance field to corresponding time steps and render the images using volume rendering based on the warped radiance field. Fig. 2 shows the overview of our method. In the following, we will introduce three main components of our method: voxel grid based canonical field, differentiable forward warping, and volume rendering. Finally, the model optimization is represented, including the proposed loss functions and training strategy.

3.1. Voxel Grid Based Canonical Field

Canonical Radiance Field To warp a static scene in canonical time to other time steps using forward warping, the radiance field should be represented by finite 3D points. The original heavy MLP in NeRF [26] is not practical, as we cannot query infinite 3D points in the canonical space. Inspired by recent works [27, 44, 61] on voxel grid presentations, we propose to use a learnable voxel radiance feature grid $\mathbf{V}_{\text{Rf}}^{\text{Can}}$ and a lightweight MLP network F_{θ_1} to model the radiance of the static scene as shown in Fig. 2. The canonical radiance field is defined as

$$\mathbf{V}_{\text{R}}^{\text{Can}} = F_{\theta_1}(\mathbf{V}_{\text{p}}^{\text{Can}}, \mathbf{V}_{\text{Rf}}^{\text{Can}}), \quad (1)$$

where $\mathbf{V}_{\text{R}}^{\text{Can}} = \{\mathbf{V}_{\sigma}^{\text{Can}}, \mathbf{V}_{\text{cf}}^{\text{Can}}\}$ consists of density voxel grid $\mathbf{V}_{\sigma}^{\text{Can}}$ and color feature voxel grid $\mathbf{V}_{\text{cf}}^{\text{Can}}$. $\mathbf{V}_{\text{p}}^{\text{Can}}$ are the 3D coordinates of the voxel grid in canonical space, embedded according to NeRF [26].

Canonical Trajectory Field We propose to use Discrete Cosine Transform (DCT) [51] to represent the trajectory of

a 3D point, which ensures the smoothness of the motion. Similar to canonical radiance field, we also use a deformation feature grid $\mathbf{V}_{\text{Df}}^{\text{Can}}$ and a lightweight MLP network F_{θ_2} to model the canonical trajectory field, which is defined as

$$\mathbf{V}_{\text{T}}^{\text{Can}} = F_{\theta_2}(\mathbf{V}_{\text{p}}^{\text{Can}}, \mathbf{V}_{\text{Df}}^{\text{Can}}), \quad (2)$$

where $\mathbf{V}_{\text{T}}^{\text{Can}}$ contains the DCT coefficients of the voxels. Given time step t , we could get the deformation flow of these voxels from canonical to time step t by

$$\mathbf{V}_{\text{flow}}^t = f_{\text{DCT}^{-1}}(\mathbf{V}_{\text{T}}^{\text{Can}}, t) - f_{\text{DCT}^{-1}}(\mathbf{V}_{\text{T}}^{\text{Can}}, \text{Can}), \quad (3)$$

where $f_{\text{DCT}^{-1}}$ is the inverse DCT transform of the DCT coefficients as described in TrajectoryNeRF [51].

3.2. Differentiable Forward Warping

Forward warping has two major issues: many-to-one and one-to-many. Fig. 3 shows a simple example, where we aim to warp the source to the target using forward warping: a) Many-to-one: if a circle shrinks, there will be multiple positions corresponding to one position from the source to the target. b) One-to-many: if a circle expands, multiple positions in the target will correspond to the same position in the source, which leaves holes of the target if we warp the source to target naively. For the many-to-one issue, we propose to use the average splatting to fuse multiple values into one. For the one-to-many issue, we use an inpaint network to inpaint the missing positions of the warped grids. More details are in the following sections.

Average Splatting We propose to fuse possible multiple values from the source grid that are mapped into the same voxel position of the target grid, motivated by Softmax-splatting [31]. Specifically, we propose a simple yet effective method that calculates the ‘‘average’’ of these values with a trilinear kernel. Formally, suppose that we need to warp the source grid \mathbf{V}^{S} to target grid \mathbf{V}^{T} by the flow $f_{\text{S} \rightarrow \text{T}}$, and \mathbf{p}, \mathbf{q} are indexes of a voxel grid. We define $\mathbf{V}^{\text{T}} = F_{\text{warp}}(\mathbf{V}^{\text{S}}, f_{\text{S} \rightarrow \text{T}})$ as follows,

$$\mathbf{V}^{\text{T}}[\mathbf{p}] = \frac{\sum_{\mathbf{q} \in \mathbf{V}^{\text{S}}} b[\mathbf{u}] \cdot \mathbf{V}^{\text{S}}[\mathbf{q}]}{\sum_{\mathbf{q} \in \mathbf{V}^{\text{S}}} b[\mathbf{u}]}, \quad (4)$$

$$b[\mathbf{u}] = \prod \max(0, 1 - |\mathbf{u}_i|), i \in \{x, y, z\}, \quad (5)$$

$$\mathbf{u} = (\mathbf{q} + f_{\text{S} \rightarrow \text{T}}[\mathbf{q}]) - \mathbf{p}, \quad (6)$$

where x, y, z are three axes of the voxel grid, and $\mathbf{u}_i \in \mathbb{R}$ is one element of the vector $\mathbf{u} \in \mathbb{R}^3$.

So, we warp the canonical radiance field to time step t by

$$\mathbf{V}_{\text{R}}^t = F_{\text{warp}}(\mathbf{V}_{\text{R}}^{\text{Can}}, \mathbf{V}_{\text{flow}}^t). \quad (7)$$

Inpaint Network Another issue of forward warping is one-to-many, which leaves holes in the voxel grid after

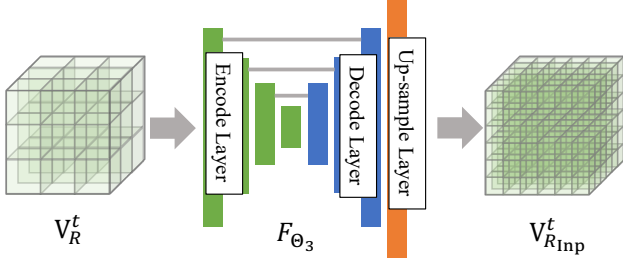


Figure 4. **Inpaint network.** It consists of a 3D U-Net structure and an up-sampling layer, which could inpaint and upsample the input voxel grid.

warping. To resolve this issue, we propose building an inpaint network to fill these holes. As shown in Fig. 4, we modify a 3D U-Net structure network [53] followed by an up-sampling layer, which is $\mathbf{V}_{R_{Inp}}^t = F_{\theta_3}(\mathbf{V}_R^t)$, where $\mathbf{V}_{R_{Inp}}^t$ is the inpainted and up-sampled voxel grid.

This structure helps encoding layers to learn information from local neighboring voxels and decoding layers to recover the original resolution with filled content. We only feed the warped voxel grid to the inpaint network, which has no access to temporal information. Thus, the inpaint network can not learn to warp, limiting its ability only to inpainting. Since warping operations and the inpaint network are relatively expensive to compute, we propose to attach an up-sample layer after the U-Net structure. In this way, we could warp the grids with a low resolution to save time and memory, while rendering with a higher resolution for better image quality. The up-sample layer consists of an interpolation layer, convolution layers, and activation layers.

3.3. Volume Rendering

After we get the radiance voxel grid at time t , the pixel colors of the image rays can be rendered using volume rendering techniques [26]. Given a ray $\mathbf{r}(w) = \mathbf{o} + w\mathbf{d}$ emitted from the camera center \mathbf{o} with view direction \mathbf{d} through a given pixel on the image plane, we render the corresponding pixel color $\mathbf{C}_{Inp}(\mathbf{r}) = F_{render}(\mathbf{V}_{R_{Inp}}^t, \mathbf{r})$. To do this, we get all 3D points \mathbf{p} that the ray intersects with voxel grids. Next, trilinear interpolation is applied to obtain the density and color feature of each 3D point.

$$(\sigma, \mathbf{c}_f) = F_{inter}(\mathbf{V}_{R_{Inp}}^t, \mathbf{p}). \quad (8)$$

Then, we get the color of this 3D point \mathbf{p} by

$$\mathbf{c} = \begin{cases} \mathbf{c}_f & \text{if } \mathbf{c}_f \in \mathbb{R}^3, \\ F_{\theta_4}(\mathbf{c}_f, \mathbf{d}) & \text{otherwise,} \end{cases} \quad (9)$$

where F_{θ_4} is a small MLP network that is used to produce colors that depend on the direction of the ray.

Finally, the pixel color could be rendered by

$$\mathbf{C}(\mathbf{r}) = \sum_{k=1}^K T(w_k) \alpha(\sigma(w_k)\delta_k) \mathbf{c}(w_k), \quad (10)$$

$$T(w_k) = \exp\left(-\sum_{j=1}^{k-1} \sigma(w_j)\delta_j\right), \quad (11)$$

where K is the number of sampled points along the ray, δ_k is the distance between adjacent samples along the ray, and $\alpha(\sigma(w_k)\delta_k) = 1 - \exp(-\sigma(w_k)\delta_k)$.

Note that the warped radiance grid \mathbf{V}_R^t should also render reasonable images. So we directly up-sample \mathbf{V}_R^t to get $\mathbf{V}_{R_{Up}}^t$ without inpainting, and render the corresponding pixel color $\mathbf{C}_{Up}(\mathbf{r}) = F_{render}(\mathbf{V}_{R_{Up}}^t, \mathbf{r})$.

3.4. Model Optimization

To optimize the model, we design a series of loss functions. **Photometric Loss** First, the most important loss is the photometric loss, which is the mean square error (MSE) of the rendered color and ground truth color $\mathbf{C}_{gt}(\mathbf{r})$.

$$\begin{aligned} \mathcal{L}^{photo} = & \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \|\mathbf{C}_{Inp}(\mathbf{r}) - \mathbf{C}_{gt}(\mathbf{r})\|_2^2 \\ & + \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \|\mathbf{C}_{Up}(\mathbf{r}) - \mathbf{C}_{gt}(\mathbf{r})\|_2^2, \end{aligned} \quad (12)$$

where \mathcal{R} is the set of rays sampled in one batch.

Following DVGO [44], we use \mathcal{L}^{pic} to directly supervise the color of sampled points. The intuition is that sampled points with larger weights contribute more to the rendered color. Also, we use the background entropy loss \mathcal{L}^{bg} to encourage the densities concentrating on either the foreground or the background.

Inpaint Network Loss As the inpaint network is used to complete the warped voxel grids, the output of the inpaint network should be close to the input. Therefore, we propose to use the L1 norm of the difference between the input and output of the inpaint network.

$$\mathcal{L}^{vdiff} = \left\| \mathbf{V}_{R_{Inp}}^t - \mathbf{V}_{R_{Up}}^t \right\|_1. \quad (13)$$

Regularization Terms Since modeling the appearance and motion of a dynamic scene from monocular images is a non-trivial problem, we propose a series of regularization terms.

First, we encourage most 3D points in canonical space to be static by introducing \mathcal{L}^{flow} to be the L1 norm of \mathbf{V}_{flow}^t . Second, denoting \mathcal{L}^{tv} as the total variation function. We compute $\mathcal{L}^{tv}(\mathbf{V}_{\sigma}^{Can})$ to ensure the spatial smoothness of the density in canonical space, and introduce $\mathcal{L}^{tv}(\mathbf{V}_{flow}^t)$ to encourage the spatial smoothness of the motions. Finally, inspired by RegNeRF [29], we render depths \mathbf{D} of image patches sampled from random views and minimize their total variations $\mathcal{L}^{tv}(\mathbf{D})$.

Table 1. **Quantitative comparison on D-NeRF Dataset.** Comparison of our method with others on LPIPS and PSNR/SSIM. The **Red** text indicates the best and **blue** text is the second best result.

Methods	Type	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
T-NeRF [36]	N	29.51	0.95	0.08
TiNeuVox-S (100 ³) [10]	NPC	30.75	0.96	0.07
TiNeuVox-B (160 ³) [10]	NPC	32.67	0.97	0.04
D-NeRF [36]	PC	30.50	0.95	0.07
NDVG (160 ³) [15]	PC	30.54	0.96	0.05
Ours (80 ³)	PC	32.68	0.97	0.04

Overall Loss The overall loss function used for optimization defined as follow, where $w_{1\sim7}$ are weights to balance each component in the final coarse loss.

$$\mathcal{L} = \mathcal{L}^{\text{photo}} + w_1 \mathcal{L}^{\text{ptc}} + w_2 \mathcal{L}^{\text{bg}} + w_3 \mathcal{L}^{\text{flow}} + w_4 \mathcal{L}^{\text{vdiff}} + w_5 \mathcal{L}^{\text{tv}}(\mathbf{V}_\sigma^{\text{Can}}) + w_6 \mathcal{L}^{\text{tv}}(\mathbf{V}_{\text{flow}}^t) + w_7 \mathcal{L}^{\text{tv}}(\mathbf{D}), \quad (14)$$

Training Strategy First, we propose *progressive training*, which starts training with images close to the canonical time and progressively adds images with farther time steps until all images are added. Second, we set up a *coarse-to-fine training strategy*. In the coarse stage, we set the color feature to be the color itself, without considering the ray directions. We compute a smaller bounding box with the proxy geometry learned from the coarse stage, which could filter a large portion of empty space for the fine stage training. The model trained in the fine stage is our final model, and we set the color features to be high dimensional features and model the ray direction dependency with F_{θ_4} .

4. Experiments

4.1. Baselines and Evaluation Datasets

Baselines The methods compared in this paper are classified into three types. First, we compare methods which are *non-canonical based (N)*, including NeRF[26], and T-NeRF[36]. Second, *physical canonical based methods (PC)* contain D-NeRF [36], NDVG[15] and Ours. These methods set their canonical space as one frame of the whole timeline by explicitly giving zeros to estimated flows at the canonical time. The canonical space of these methods should have reasonable physical 3D reconstructions of the scene at the canonical time. Third, *non-physical canonical based methods (NPC)* consists of NV[23], Nerfies[32], HyperNeRF[33] and TiNeuVox[10]. Their canonical space geometries do not necessarily have physical meanings.

D-NeRF Dataset We evaluate our method on 8 dynamic scenes of D-NeRF [36] dataset. We report several common metrics for the evaluation: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM) and Learned Perceptual Image Patch Similarity (LPIPS) [65].




Figure 5. **Qualitative comparison on D-NeRF Dataset.** We show some novel view synthesized images on the selected test set of the dataset. Comparing ours with ground truth, D-NeRF [36] and TiNeuVox [10]. Our model yields cleaner images with more details.

HyperNeRF Dataset Besides synthetic datasets, we also conduct experiments on real scenes, proposed by HyperNeRF [33]. This dataset uses a multi-view camera rig consisting of 2 phones to capture real unbounded scenes with challenging rigid and non-rigid deformations.

NHR Dataset NHR [54] dataset capture dynamic human data using a multi-camera dome system with up to 80 cameras arranged on a cylinder. We conduct experiments on four scenes where the performers are in different clothing and perform different actions. We test with 100 frames of each scene, selecting 90% views for training and 10% views for testing.

Lego Complete Dataset The D-NeRF [36] dataset only contains 20 test images for each scene, and the evaluations of D-NeRF do not contain images in the canonical space. However, for *physical canonical based methods (PC)*, canonical space geometry is important, as it affects images quality rendered at the canonical time and reflects the quality of flow estimation. Hence, we build a new dataset, named *Lego Complete Dataset*, which animates the object *LEGO* in D-NeRF dataset with 3 different motion patterns. For each scene, the test set is split into 3 categories to evaluate three abilities: *space interpolation* (rand views for each train time step), *time interpolation* (times between trained time steps), and *canonical interpolation* (rand views for canonical time step) abilities. For more details, please refer to our *supp*.

Besides the image quality, we also consider geometry precision by evaluating the depth of each test image to analyze the reconstructed geometry. For depth metrics, we use the relative error $\delta = \max\left(\frac{d}{d_{\text{gt}}}, \frac{d_{\text{gt}}}{d}\right)$, where d is the estimated depth and d_{gt} is the ground truth depth. We report the percentage of the pixels with $\delta < \text{Threshold}$.

Table 2. **Quantitative comparison on real scene HyperNeRF dataset.** Comparison of our method with others on PSNR and MS-SSIM.

Methods	Type	3D Printer		Broom		Chicken		Peel Banana		Mean	
		PSNR↑	MS-SSIM↑	PSNR↑	MS-SSIM↑	PSNR↑	MS-SSIM↑	PSNR↑	MS-SSIM↑	PSNR↑	MS-SSIM↑
NeRF [26]	N	20.7	0.780	19.9	0.653	19.9	0.777	20.0	0.769	20.1	0.745
NV [23]	NPC	16.2	0.665	17.7	0.623	17.6	0.615	15.9	0.380	16.9	0.571
Nerfies [32]	NPC	20.6	0.830	19.2	0.567	26.7	0.943	22.4	0.872	22.2	0.803
HyperNeRF [33]	NPC	20.0	0.821	19.3	0.591	26.9	0.948	23.3	0.896	22.4	0.814
TiNeuVox-S (100 ³) [10]	NPC	22.7	0.836	21.9	0.707	27.0	0.929	22.1	0.780	23.4	0.813
TiNeuVox-B (160 ³) [10]	NPC	22.8	0.841	21.5	0.686	28.3	0.947	24.4	0.873	24.3	0.837
NDVG (160 ³) [15]	PC	22.4	0.839	21.5	0.703	27.1	0.939	22.8	0.828	23.3	0.823
Ours (70 ³)	PC	22.8	0.845	21.9	0.715	28.0	0.944	24.3	0.865	24.2	0.842




Figure 6. **Qualitative comparison on HyperNeRF Dataset.** Our results are closer to ground truth than other methods.

4.2. Experimental Results

Evaluation on D-NeRF Dataset We compared our method with the canonical-based backward flow methods D-NeRF [36], NDVG [15] and TiNeuVox [10] on D-NeRF Dataset. As shown in Table 1, our method achieves significant improvements in terms of all three metrics for *physical canonical based methods*. On average, our method improves PSNR by 2.18 compared with D-NeRF and 2.14 compared with NDVG. For *non-physical canonical based methods*, our method also achieves the best result. Ours show obviously better qualitative rendering quality compared with TiNeuVox-S (100³) [10] with even smaller voxel resolution. More detailed statistics are provided in *supp*.

We also provide some visual comparisons in Fig. 5. Ours could recover accurate and detailed images, *e.g.*, the helmet and arm in the top scene, and could also produce cleaner boundaries, *e.g.* the hand and feet in the bottom scene.

Evaluation on HyperNeRF Dataset We further compare our method with some highly related works on real scene dataset proposed by [33]. As shown in Table 2, our method achieves consistently better performance among *physical canonical based methods*. Our method has a voxel grid resolution limitation, due to we need to warp the whole voxel grid for rendering. Compared with other voxel grid based methods, our voxel grid resolution is limited to 70³, along with a 160³ static voxel grid for background modeling. However, the performance of our method is still com-



Figure 7. **Visual comparisons on NHR dataset.**

parable with SOTA of *non-physical canonical based methods*, proving the effectiveness of our forward flow design.

In Fig. 6, we show visual comparisons with NDVG [15], and TiNeuVox [10]. The forward flow helps to recover the correct structure of dynamic objects, like the eye and edges of the chicken (top left) and the broom (top right). The relatively low resolution may harm the ability to recover very fine details, like the patterns on the 3D printer compared with TiNeuVox-B [10] (bottom right).

Evaluation on NHR Dataset We also test our method on NHR dataset. Quantitative results are shown in Table 3, which shows ours achieving best results consistently. We show qualitative comparison in Fig. 7, and our method could render clean and detailed images.

Table 3. **Quantitative comparison on NHR dataset.** The **red** text indicates the best and **blue** text is the second best result.

Methods	type	Sport 1			Sport 2			Sport 3			Basketball			Mean		
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
TiNeuVox-S [10]	NPC	26.06	0.93	0.10	25.98	0.93	0.11	25.90	0.93	0.11	23.75	0.91	0.14	25.42	0.92	0.12
TiNeuVox-B [10]	NPC	26.44	0.93	0.10	26.68	0.94	0.10	26.09	0.93	0.11	25.06	0.92	0.12	26.07	0.93	0.11
NDVG [15]	PC	23.66	0.89	0.15	24.43	0.91	0.13	22.54	0.88	0.16	22.55	0.89	0.17	23.29	0.89	0.15
Ours	PC	27.71	0.95	0.08	27.89	0.95	0.08	27.57	0.94	0.08	24.85	0.93	0.11	27.00	0.94	0.09

Table 4. **Quantitative comparison on Lego Complete Dataset.** We report the average PSNR, LPIPS and the average relative depth error.

Methods	Type	Space Interpolation				Time Interpolation				Canonical Interpolation			
		PSNR \uparrow	LPIPS \downarrow	$\delta < 1.25$	$\delta < 1.25^2$	PSNR \uparrow	LPIPS \downarrow	$\delta < 1.25$	$\delta < 1.25^2$	PSNR \uparrow	LPIPS \downarrow	$\delta < 1.25$	$\delta < 1.25^2$
D-NeRF [36]	PC	23.98	0.15	98.720	99.328	24.18	0.15	98.854	99.378	17.36	0.23	97.341	98.868
NDVG (160 ³) [15]	PC	28.12	0.06	99.655	99.853	28.19	0.06	99.749	99.953	19.90	0.11	94.828	97.560
Ours_notv (80 ³) ¹	PC	27.58	0.07	99.758	99.958	27.63	0.07	99.788	99.966	23.56	0.13	99.238	99.820
Ours_noup (80 ³) ²	PC	26.11	0.09	99.596	99.981	26.19	0.09	99.636	99.983	24.64	0.10	99.230	99.925
Ours_noinp (80 ³) ³	PC	27.21	0.08	99.783	99.979	27.27	0.08	99.797	99.980	23.91	0.09	99.586	99.963
Ours (80 ³)	PC	28.18	0.06	99.818	99.985	28.22	0.06	99.839	99.987	25.63	0.07	99.654	99.968

¹ not use the all total variation losses ² not up-sample the voxel grid ³ not use the inpaint network

4.3. Method Analysis

Flow Estimation and Canonical Space Geometry Since evaluating the estimated flow with ground truth is not practical, we choose to compare the reconstructed canonical geometry of *physical canonical based methods*. The idea is that if the canonical space is one frame of the whole timeline (like D-NeRF [36] and NDVG [15]), better flow estimation will result in better canonical scene geometry. We test these methods on Lego Complete Dataset, and all set the canonical space as the first frame of the image sequence.

In Table 4, D-NeRF can not reconstruct the canonical radiance field well, since both the image and depth quality in canonical interpolation are significantly worse than those in space interpolation and time interpolation. NDVG [15] could recover much better images and geometry for other time steps, but the differences between canonical interpolation with the other two are even more significant. This means NDVG [15] has difficulty estimating correct backward flows. To render good quality images at other time steps, the correct geometries (depth) at these time steps are warped back into distorted canonical geometry due to incorrect backward flows.

This is the proof of our claim that the non-smooth and discontinuous nature of the backward flow deformation field makes it difficult to fit with smooth functions, especially for NDVG [15], which has limited MLP capacity for fast training. However, our method performs significantly better for both image quality and depth, and there are fewer variations between canonical, space, and temporal interpolations. This demonstrates the benefit of forward warping, which makes it easier for the deformation model to learn deformation flows from canonical time to other times.

We compare the canonical image of our method with D-NeRF [36] in Fig. 8. Our method could recover the correct

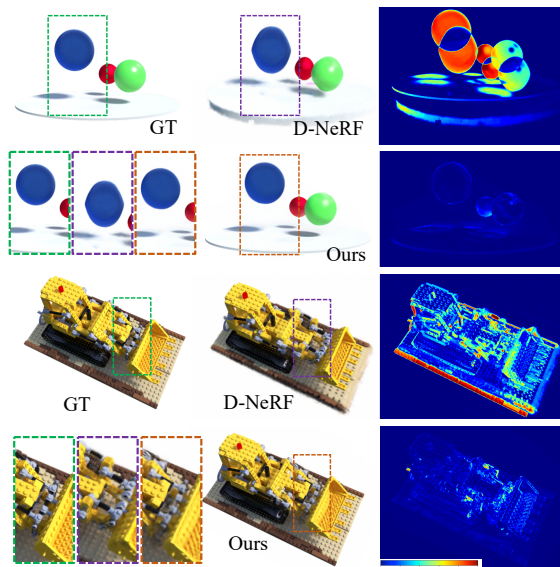


Figure 8. **Canonical qualitative comparison.** We show canonical radiance field comparison with D-NeRF [36]. Given the error map between the ground truth and rendered images, we can see that the canonical frame yielded by ours is closer to the ground truth. The results of D-NeRF are blurry and have large displacements.

geometry and image details of the canonical frame. For example, the arms of the bucket produced by D-NeRF are at the “mean” positions of the whole trajectory, and ours recover the correct status.

Trajectory Visualization We show the learned DCT trajectories generated by our canonical trajectory field in Fig. 7 and Fig. 9. The canonical trajectory field is capable of recovering reasonable object point trajectories. Backward warping designs cannot track the motions of the same object point, so this is not a viable solution. With this feature, geometry constraints, motion models, and prior knowledge could be introduced in future work.




Figure 9. **Trajectory learned by the canonical trajectory field.** Light blue is the canonical frame, the curve represents the historical motion trajectory. More videos can be found in *supp*.

Regularization terms We study the effect of all regularization terms we proposed, including $\mathcal{L}^{\text{flow}}$, $\mathcal{L}^{\text{vdiff}}$ and \mathcal{L}^{tv} in Table 5. We could observe that all these three regularization terms has positive effect on the performance, but the improvement is minor. This proves the improvement of our method compared with others come from the forward warping design we proposed in the paper. Also, backward flow based method NDVG [15] use similar regularization terms with ours, and our method has clear advantage compared with NDVG [15].

Canonical setup In our method, we set canonical time to be the first frame for D-NeRF dataset to compare with *physical canonical based method* and the middle frame for HyperNeRF dataset for better performance. Setting canonical time to be the middle frame helps improving the performance as the state of the middle frame geometry is closer to other time steps compared with the first frame. To prove this, we set canonical time to be middle frame in Table 5 (can-time t at mid), and the PSNR improves slightly.

Photometric loss for \mathbf{V}_{RUp} We use photometric terms on \mathbf{V}_{RUp} to make sure the warped grid before inpainting could already render reasonable images. This make sure the UNet is actually doing ‘inpainting’. Also, we do observe inpainting and upsampling could refine grids. In Fig. 7 and Fig. 9, the trajectories are reasonable which means we do learn trajectories without inpainting overfitting. We also test w/o photometric terms of \mathbf{V}_{RUp} in Table 5 (w/o \mathbf{V}_{RUp} photo) and the performance drops as there is no direct supervise signal for trajectory learning.

More Ablations We study the effects of the total variation losses, the up-sample strategy, and the inpaint network we proposed in Table 6. The total variation losses help to

Table 5. **Ablations.** Mean of Hell Warrior, Mutant, Hook and Bouncing Balls in D-NeRF dataset.

Method	Ours	w/o \mathcal{L}^{tv}	w/o $\mathcal{L}^{\text{flow}}$	w/o $\mathcal{L}^{\text{vdiff}}$	can-time t at mid	w/o \mathbf{V}_{RUp} photo
PSNR	33.75	33.30	33.74	33.57	34.09	31.56
SSIM	0.979	0.974	0.978	0.977	0.980	0.966
LPIPS	0.040	0.048	0.039	0.041	0.037	0.057

Table 6. **Ablation of losses on two dataset.**

Methods	D-NeRF		HyperNeRF	
	PSNR \uparrow	LPIPS \downarrow	PSNR \uparrow	MS-SSIM \uparrow
Ours_notv ¹	33.09	0.046	23.84	0.831
Ours_noup ²	32.03	0.043	23.72	0.822
Ours_noinp ³	31.78	0.044	23.25	0.814
Ours	32.68	0.037	24.23	0.842

improve visual quality (LPIPS) in D-NeRF dataset as it applies smoothness over the deformation and radiance space, but may harm the high-frequency details (PSNR) for some of the scenes in D-NeRF dataset. For HyperNeRF dataset, the total variation losses improve the results, as the camera settings are more challenging than D-NeRF dataset. However, this improvement is minor compared with the other two factors. Since the resolution is critical for voxel grid based method, up-sampling operation of inpaint network is vital to improve the performance. Finally, the inpaint network plays the most essential role among these three factors in our method. More ablations are provided in our *supp*.

5. Conclusion

This paper presents a canonical-based representation with forward warping for novel view synthesis of dynamic scenes. Our method models a static scene at the canonical field and forward warp this whole field to other time steps for dynamic scene rendering. To address the many-to-one and one-to-many mapping difficulties, we present a differentiable forward warping with an average splatting process and an inpaint network. Our proposed forward warping pipeline achieves SOTA performance on the public datasets and our newly built dataset, especially for canonical frames.

Limitations and Future Directions Our current implementation is relatively memory-consuming, especially for real scenes. Also, the training speed is relatively slow (one day for each scene). Since we have a smooth trajectory field thanks to the forward warping, additional constraints and motion models could be introduced to learn better trajectories in future works.

Acknowledgements This research was supported in part by the National Natural Science Foundation of China (Grant Nos. 62271410, 62293482, 62202409), the Fundamental Research Funds for the Central Universities, and Shenzhen Science and Technology Program (Grant No. RCBS20221008093241052).

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] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerf: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [4] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001. 2
- [5] 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
- [6] Shenchang Eric Chen and Lance Williams. View interpolation for image synthesis. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, 1993. 2
- [7] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [8] 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
- [9] Yilun Du, Yinan Zhang, Hong-Xing Yu, Joshua B Tenenbaum, and Jiajun Wu. Neural radiance flow for 4d view synthesis and video processing. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 1, 3
- [10] 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. *ACM SIGGRAPH Asia*, 2022. 3, 6, 7, 8
- [11] 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
- [12] Wanshui Gan, Hongbin Xu, Yi Huang, Shifeng Chen, and Naoto Yokoya. V4d: Voxel for 4d novel view synthesis. *arXiv preprint arXiv:2205.14332*, 2022. 3
- [13] 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. 1, 2
- [14] Ned Greene. Environment mapping and other applications of world projections. *IEEE Computer Graphics and Applications*, 1986. 2
- [15] 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, 3, 6, 7, 8, 9
- [16] 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
- [17] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi. Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 2016. 2
- [18] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996. 2
- [19] 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. 1, 3
- [20] 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
- [21] Jia-Wei Liu, Yan-Pei Cao, Weijia Mao, Wenqiao Zhang, David Junhao Zhang, Jussi Keppo, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Devrf: Fast deformable voxel radiance fields for dynamic scenes. *arXiv preprint arXiv:2205.15723*, 2022. 3
- [22] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [23] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Transactions on Graphics (TOG)*, 38(4):65:1–65:14, July 2019. 6, 7
- [24] 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
- [25] 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
- [26] 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, 6, 7

- [27] 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, 4
- [28] 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*, 2021. 2
- [29] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 5
- [30] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [31] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 4
- [32] 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, 6, 7
- [33] 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. 2, 3, 6, 7
- [34] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)*, 2017. 2
- [35] Martin Píala and Ronald Clark. Terminerf: Ray termination prediction for efficient neural rendering. In *Proceedings of the International Conference on 3D Vision (3DV)*, 2021. 2
- [36] 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, 7, 8
- [37] 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
- [38] 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
- [39] 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
- [40] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [41] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [42] Harry Shum and Sing Bing Kang. Review of image-based rendering techniques. In *Visual Communications and Image Processing*, 2000. 2
- [43] 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
- [44] 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, 4, 5
- [45] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [46] 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*, 2020. 2
- [47] 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 *ACM SIGGRAPH 2021*, 2021. 2
- [48] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 2019. 2
- [49] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2021. 1, 3
- [50] 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
- [51] Chaoyang Wang, Ben Eckart, Simon Lucey, and Orazio Gallo. Neural trajectory fields for dynamic novel view synthesis. *arXiv preprint arXiv:2105.05994*, 2021. 1, 3, 4
- [52] 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

- [53] Adrian Wolny, Lorenzo Cerrone, Athul Vijayan, Rachele Tofanelli, Amaya Vilches Barro, Marion Louveaux, Christian Wenzl, Sören Strauss, David Wilson-Sánchez, Rena Lymbouridou, Susanne S Steigleder, Constantin Pape, Alberto Bailoni, Salva Duran-Nebreda, George W Bassel, Jan U Lohmann, Miltos Tsiantis, Fred A Hamprecht, Kay Schneitz, Alexis Maizel, and Anna Kreshuk. Accurate and versatile 3d segmentation of plant tissues at cellular resolution. *eLife*, 9:e57613, 2020. 5
- [54] Minye Wu, Yuehao Wang, Qiang Hu, and Jingyi Yu. Multi-view neural human rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 6
- [55] 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. 1, 3
- [56] 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
- [57] 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
- [58] 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. *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2
- [59] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Ronen Basri, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 1
- [60] 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
- [61] 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, 4
- [62] 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
- [63] 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
- [64] Kai Zhang, Gernot Riegler, Noah Snaveley, and Vladlen Koltun. NERF++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 2
- [65] 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
- [66] 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