

Neural Deformable Voxel Grid for Fast Optimization of Dynamic View Synthesis

Xiang Guo^{1*}, Guanying Chen^{2*}, Yuchao Dai^{1†}, Xiaoqing Ye³,
Jiada Sun¹, Xiao Tan³, and Errui Ding³

¹ Northwestern Polytechnical University

² FNii and SSE, CUHK-Shenzhen

³ Baidu Inc.

{guoxiang,sunjiadai}@mail.nwpu.edu.cn, chenguanying@cuhk.edu.cn,
daiyuchao@nwpu.edu.cn, {yexiaoqing,dingerrui}@baidu.com,
tanxchong@gmail.com

Abstract. Recently, Neural Radiance Fields (NeRF) is revolutionizing the task of novel view synthesis (NVS) for its superior performance. In this paper, we propose to synthesize dynamic scenes. Extending the methods for static scenes to dynamic scenes is not straightforward as both the scene geometry and appearance change over time, especially under monocular setup. Also, the existing dynamic NeRF methods generally require a lengthy per-scene training procedure, where multi-layer perceptrons (MLP) are fitted to model both motions and radiance. In this paper, built on top of the recent advances in voxel-grid optimization, we propose a fast deformable radiance field method to handle dynamic scenes. Our method consists of two modules. The first module adopts a deformation grid to store 3D dynamic features, and a light-weight MLP for decoding the deformation that maps a 3D point in the observation space to the canonical space using the interpolated features. The second module contains a density and a color grid to model the geometry and density of the scene. The occlusion is explicitly modeled to further improve the rendering quality. Experimental results show that our method achieves comparable performance to D-NeRF using only 20 minutes for training, which is more than 70× faster than D-NeRF, clearly demonstrating the efficiency of our proposed method.

Keywords: Dynamic View Synthesis · Neural Radiance Fields · Voxel-grid Representation · Fast Optimization.

1 Introduction

Novel view synthesis (NVS) is a long-standing problem in computer vision and graphics, and has many applications in augmented reality, virtual reality, content creation, *etc.* Recently, neural rendering methods have achieved significant progress in this problem [31,60,34]. In particular, the neural radiance fields

* Authors contributed equally to this work. † Yuchao Dai is the corresponding author.

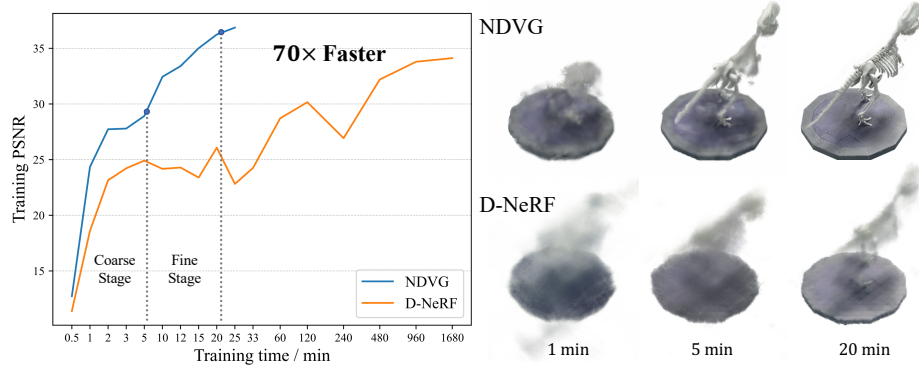


Fig. 1. **Neural Deformable Voxel Grid (NDVG)** for fast optimization of dynamic view synthesis. Left side of the figure shows that our method achieves a super fast convergence in 20 minutes, which is 70× faster than the D-NeRF method. Right side of the figure visualizes the results after training with 1, 5 and 20 minutes.

(NeRF) [31] produces photorealistic rendering by representing a static scene with a multi-layer perception (MLP), which maps a 5D input (3D coordinate and 2D view direction) to its density and color. Recently, a series of works extend NeRF based framework from static scenes to dynamic scenes [13,22,58,52,35,39,55,8].

Novel view synthesis of a dynamic scene from a monocular video is still a very challenging problem. Besides the difficulties to recover motions and geometries with only one observation at each time step, the training process usually takes days which hinders applications in practice. The NeRF based methods, at each iteration, require millions of network queries to obtain colors and densities of the sampled points for the sampled rays, based on which volume rendering computes the pixel colors [18]. In dynamic condition, the methods are even more complex with deformation model, *e.g.*, D-NeRF [39] optimizes a large deformation network and a canonical network to fit a dynamic scene, requiring more than 27 hours to converge. How to develop an efficient and accurate dynamic view synthesis method remains an open problem.

In a static scenario, to reduce the training time for a scene, some methods propose first to train the model on a dataset consisting of multiple scenes [5,65,53,56], and then finetune it on the target scene, reducing the optimization time to several minutes. However, these methods rely on a large training dataset and a lengthy pre-training time.

Very recently, the voxel-grid representation has been exploited to speed up the optimization of radiance fields [64,48,32]. These methods are able to optimize a scene representation from scratch within just a few minutes, significantly accelerating the training speed without any pre-training. The key idea is to replace the time-consuming deep network query with the fast trilinear voxel-grid interpolation. However, these methods are tailored for static scenes and cannot be directly applied to handle dynamic scenes.

To the best of our knowledge, there are few works for fast optimization dynamic NeRF. Also, it is challenging to design a compact deformation model, while having enough capacity. In this paper, we propose a fast optimization method, named NDVG for dynamic scene view synthesis based on the voxel-grid representation. Our method consists of a deformation module and a canonical module. The deformation module maps a 3D point in the observation space to canonical space, and volume rendering is performed in the canonical space to render the pixel color and compute image reconstruction loss. In contrast to static scene applications where the occlusion is only determined by the viewing direction when the scene is known, the occlusion in the dynamic scene NeRF is determined by both view direction and motion (or view time) and should be taken well care of. Hence, we designed a particular occlusion handling module to explicitly model the occlusion to further improve the rendering quality.

In summary, the key contributions of this paper are as follows:

- We propose a fast deformable radiance field method based on the voxel-grid representation to enable space-time view synthesis for dynamic scenes. To the best of our knowledge, this is the first method that integrates the voxel-grid optimization with deformable radiance field.
- We introduce a deformation grid to store the 3D dynamic features and adopt a light-weight MLP to decode the feature to deformation. Our method explicitly models occlusion to improve the results.
- Our method produces rendering results comparable to D-NeRF [39] within only 20 minutes, which is more than 70× faster (see Fig. 1).

2 Related Work

Novel View Synthesis Rendering a scene from arbitrary views has a long history in both vision and graphics [4,6,21,14], and surveys of recent methods can be found in [44,50,49]. Traditional methods need to explicitly build a 3D model for the scene, such as point clouds [1] or meshes [42,43,51,16], and then render a novel view from this geometry. Another category of methods explicitly estimates depth and then uses it to warp pixels or learned feature to a novel view, such as [19,37,7,42,43,11,59]. Numerous other works using multi-plane images (MPIs) to represent scenes [10,30,46,47,69,54,17,15], but the MPIs representation can only support relatively limited viewpoint changes during inference.

Neural Scene Representation Recently, neural scene representations dominate novel view synthesis. In particular, Mildenhall *et al.* propose NeRF [31] to use MLPs to model a 5D radiance field, which can render impressive view synthesis for static scenes captured. Since then, many follow-up methods have extended the capabilities of NeRF, including relighting [3,45,68], handling in-the-wild scenarios [29], extending to large unbounded 360° scenes [66], removing the requirement for pose estimation [57,23], incorporating anti-aliasing for multi-scale rendering [2], and estimating the 6-DoF camera poses [61].

Fast NeRF Rendering and Optimization Rendering and optimization in NeRF-like schemes are very time-consuming, as they require multiple samples

along each ray for color accumulation. To speed up the rendering procedure, some methods predict the depth or sampling near the surface to guide more efficient sampling [33,38]. Other methods use the octree or sparse voxel grid to avoid sampling points in empty spaces [26,64,28]. In addition, some methods subdivided the 3D volume into multiple cells that can be processed more efficiently, such as DeRF [40] and KiloNeRF [41]. AutoInt [24] reduces the number of evaluations along a ray by learning partial integrals. However, these methods still need to optimize a deep implicit model, leading to a lengthy training time. To accelerate the optimization time on a test scene, some methods first train the model on a large dataset to learn a scene prior, and then finetune it on the target scene [5,53,65,56]. However, these methods require time-consuming pretraining.

More recently, the voxel-grid representation has been exploited to speed up the optimization of radiance field [63,48,32]. Plenoxels [63] represents a scene as a 3D grid with spherical harmonics, which can be optimized from calibrated images via gradient methods and regularization without any neural components. Similarly, DVGO [48] optimizes a hybrid explicit-implicit representation that consists of a dense grid and a light-weight MLP. Although these methods achieve a fast optimization speed, they are only applicable to static scenes and thus cannot be used to render dynamic scenes.

Dynamic Scene Modeling Recently, several concurrent methods have extended NeRF to deal with dynamic scenes [13,22,58,52,35,39,55,8].

NeRFlow [8] learns a 4D spatial-temporal representation of a dynamic scene from a set of RGB images. Yoon *et al.* [62] propose to use an underlying 4D reconstruction, combining single-view depth and depth from multi-view stereo to render virtual views with 3D warping. Gao *et al.* [13] jointly train a time-invariant model (static) and a time-varying model (dynamic), and regularize the dynamic NeRF by scene flow estimation, finally blending the results in an unsupervised manner. NSFF [22] models the dynamic scene as a time-variant continuous function of appearance, geometry, and 3D scene motion. DCT-NeRF [55] uses the Discrete Cosine Transform (DCT) to capture the dynamic motion and learn smooth and stable trajectories over time for each point in space.

D-NeRF [39], NR-NeRF [52] and Nerfies [35] first learn a static canonical radiance field for capturing geometry and appearance, and then learn the deformation/displacement field of the scene at each time instant w.r.t. the canonical space. Xian *et al.* [58] represent a 4D space-time irradiance field as a function that maps a spatial-temporal location to the emitted color and volume density.

Although promising results have been shown for dynamic view synthesis, these methods all require a long optimization time to fit a dynamic scene, limiting their wider applications.

As the concurrent research of our work, there are few works which aim to speed up training of dynamic NeRF [9,12,25]. TiNeuVox [9] uses a small MLP to model the deformation and uses the multi-distance interpolation to get the feature for radiance network which estimates the density and color. Compared to TiNeuVox [9], we propose a deformation feature grid to enhance the capability of the small deformation network and not effect training speed. 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. While V4D [12] focuses mainly on improving image quality, the speed up of training is not significant compared with TiNeuVox [9] and ours. DeVRF [25] 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 eases the learning of motion and geometry compared with ours which uses monocular images.

3 Neural Deformable Voxel Grid

Given an image sequence $\{I_t\}_{t=1}^T$ with camera poses $\{\mathbf{T}_t\}_{t=1}^T$ of a dynamic scene captured by a monocular camera, our goal is to develop a fast optimization method based on the radiance fields to represent this dynamic scene and support novel view synthesis at different times.

3.1 Overview

Mathematically, given a query 3D point \mathbf{p} , the view direction \mathbf{d} , and a time instance t , we need to estimate the corresponding density σ and color \mathbf{c} .

A straightforward way is to directly use an MLP to learn the mapping from $(\mathbf{p}, \mathbf{d}, t)$ to (σ, \mathbf{c}) . However, existing dynamic methods show that such a high-dimensional mapping is difficult to learn, and they propose a framework based on the canonical scene to ease learning difficulty [39, 52, 36]. These methods adopt a deformation MLP to map a 3D point in the observation space to a static canonical space as $\Psi_t : (\mathbf{p}, t) \rightarrow \Delta\mathbf{p}$. The density and color are then estimated in the canonical space as $\Psi_p : (\sigma, \mathbf{c}) = f(\mathbf{p} + \Delta\mathbf{p}, \mathbf{d})$. However, to render the network with the capability of handling complex motion, large MLPs are inevitable in existing methods and therefore result in a long optimization time (*i.e.*, from hours to days).

Motivated by the recent successes of voxel-grid optimization in accelerating the training of static radiance field [48, 63], we introduce the voxel-grid representation into the canonical scene representation based framework to enable fast dynamic scene optimization. Our method consists of a *deformation module* and a *canonical module* (see Fig. 2). The key idea is to replace most of the heavy MLP computations with the fast voxel-grid feature interpolation.

3.2 Deformation Module for Motion Modeling

Assuming the canonical space is at time t_{can} , the deformation module estimates the offset $\Delta\mathbf{p}$ of a point \mathbf{p} at any time t to the canonical space. As the input has 4 dimensions (*i.e.*, 3 for \mathbf{p} and 1 for t), it is inefficient to directly store the offsets in a 4D feature grid. Therefore, we adopt a hybrid explicit-implicit representation that consists of a 3D feature grid and a light-weight MLP to decode the interpolated feature.




Fig. 2. **Overview of our proposed method.** Our method consists of (a) a *deformation module* to model the motion of the space points and (b) a *canonical module* to model the radiance field of the static scene at the canonical time. To render a ray shooting from the camera center, we compute the deformation of all sampled points and transform the sampled points to the canonical space, where the density and color are computed. The pixel color can then be rendered by (c) volume rendering.

Neural Deformable Voxel Grid (NDVG) We first build a 3D deformation feature grid $\mathbf{G}^d \in \mathbb{R}^{N_x \times N_y \times N_z \times N_c}$ with a resolution of $N_x \times N_y \times N_z$, and the feature vector in each voxel has a dimension of N_c . For a continuous 3D coordinate \mathbf{p} , its deformation feature can be quickly queried from the deformation feature grid using trilinear interpolation:

$$\text{interp}(\mathbf{p}, \mathbf{G}^d) : (\mathbb{R}^3, \mathbb{R}^{N_x \times N_y \times N_z \times N_c}) \rightarrow \mathbb{R}^{N_c}. \quad (1)$$

To obtain the offset from the observation space to the canonical space for a query point, we introduce a light-weight MLP $F_{\theta_1}^d$. It takes the coordinate \mathbf{p} , time t , and the interpolated feature as input, and regresses the offset $\Delta \mathbf{p}$.

$$\mathbf{f}^d = \text{interp}(\mathbf{p}, \mathbf{G}^d), \quad \Delta \mathbf{p} = \begin{cases} F_{\theta_1}^d(\mathbf{p}, t, \mathbf{f}^d) & \text{if } t \neq t_{can}, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Following [31], we apply positional encoding for \mathbf{p} and t . Finally, we get the position of \mathbf{p} in the canonical space \mathbf{p}' as: $\mathbf{p}' = \mathbf{p} + \Delta \mathbf{p}$.

The deformation feature grid can provide learnable features that encode the deformation information of the points in the 3D space, so that a light-weight MLP is capable to model the deformation of the 3D space.

3.3 Canonical Module for View Synthesis

To render the pixel color of a camera ray in the observation space, we first sample K points on the ray and transform them to the canonical space via our

deformation module. The canonical module then computes the corresponding densities and colors for volume rendering to composite the pixel color.

Density and Color Grids The canonical module contains a density grid to model the scene geometry, and a hybrid explicit-implicit representation (contains a color feature grid and a light-weight MLP) to model the view-dependent effect.

The density grid $\mathbf{G}^\sigma \in \mathbb{R}^{N'_x \times N'_y \times N'_z}$ with a resolution of $N'_x \times N'_y \times N'_z$ stores the density information of the scene. Following [48], for a 3D query point \mathbf{p}' , we perform trilinear interpolation on the grid and apply post-activation with the softplus activation function to get the density: $\sigma = \text{softplus}(\text{interp}(\mathbf{p}', \mathbf{G}^\sigma))$

The color feature grid $\mathbf{G}^c \in \mathbb{R}^{H' \times W' \times D' \times N'_c}$ with a resolution of $H' \times W' \times D'$ stores the color feature of the scene, where the feature dimension is N'_c . To obtain the view dependent \mathbf{c} for a 3D point \mathbf{p}' , we use a light-weight MLP $F_{\theta_2}^c$ to decode the interpolated color feature \mathbf{f}^c . Positional encoding is applied on \mathbf{p}' and \mathbf{d} .

$$\mathbf{f}^c = \text{interp}(\mathbf{p}', \mathbf{G}^c), \quad \mathbf{c} = F_{\theta_2}^c(\mathbf{p}', \mathbf{d}, \mathbf{f}^c). \quad (3)$$

3.4 Occlusion-aware Volume Rendering

Volume rendering is the key in radiance fields based methods to render pixel color in a differentiable manner. For 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, the estimated color $\hat{\mathbf{C}}(\mathbf{r})$ of this ray is computed as

$$\hat{\mathbf{C}}(\mathbf{r}) = \sum_{k=1}^K T(w_k) \alpha(\sigma(w_k)\delta_k) \mathbf{c}(w_k), \quad T(w_k) = \exp\left(-\sum_{j=1}^{k-1} \sigma(t_j)\delta_j\right), \quad (4)$$

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

Occlusion Problem Note that for dynamic scenes, the occlusion will cause problems in the canonical scene-based methods. We assume the “empty” points (*i.e.*, non-object points) in the space are static. If a 3D point is the empty point at time t but occupied by an object point in the canonical space, the occlusion happens. This is because the ideal deformation for an empty point in the observation space is zero, then this point will be mapped to the same location (but occupied by an object point) in the canonical space (see Fig. 3 for illustration). As a result, the obtained density and color value for this empty point will be non-zero, leading to incorrect rendering color.

Occlusion Reasoning To tackle the occlusion problem, we additionally estimate an occlusion mask for a query point using the deformation MLP $F_{\theta_1}^d$. Then Eq. (2) is revised as:

$$(\Delta\mathbf{p}, w^{occ}) = F_{\theta_1}^d(\mathbf{p}, t, \mathbf{f}^D). \quad (5)$$

If a query point is an empty point at time t but is occupied by an object point in the canonical space, the estimated w^{occ} should be 0. Based on the estimated

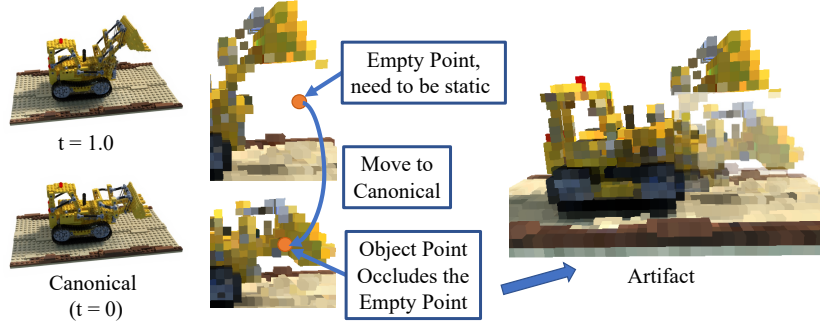


Fig. 3. **Effect of Occlusion.** If the static empty point in current time is occluded by the object point in canonical time, this empty point will bring the object point back to current time, which creates artifacts.

occlusion mask w^{occ} , we filter the estimated density and color before sending them to volume rendering to remove the influence of the occluded points:

$$\sigma' = \sigma \times w^{occ}, \quad \mathbf{c}' = \mathbf{c} \times w^{occ}. \quad (6)$$

The occlusion estimation is optimized by the image reconstruction loss in an end-to-end manner. By taking occlusion into account, our method is able to achieve better rendering quality.

4 Model Optimization

There is a trade-off between grid resolution and time consumption. Using a higher grid resolution generally leads to better image quality but at the cost of a higher computational time. Meanwhile, the space is dominated by empty space, which takes no effect on the synthesized images.

Loss Function We utilize a series of losses to optimize the model. The first is the photometric loss, which is the mean square error (MSE) between the estimated ray color $\hat{\mathbf{C}}(\mathbf{r})$ and the ground-truth color $\mathbf{C}(\mathbf{r})$ as

$$\mathcal{L}^{\text{photo}} = \frac{1}{|\mathcal{R}|} \sum_{\mathbf{r} \in \mathcal{R}} \left\| \hat{\mathbf{C}}(\mathbf{r}) - \mathbf{C}(\mathbf{r}) \right\|_2^2, \quad (7)$$

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

Also, similarly to [48], we use the sampled point color supervision \mathcal{L}^{ptc} , which optimizes the color of sampled points with top N weights that contribute most to the rendered color of the ray. In addition, we use the background entropy loss \mathcal{L}^{bg} to encourage the points to concentrate on either foreground or background.

For the deformation module, we apply the L_1 norm regularization to the estimated deformation based on the prior that most points are static:

$$\mathcal{L}^{\text{d_norm}} = \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} \sum_{i=1}^K \|\Delta \mathbf{p}_i\|_1. \quad (8)$$

Moreover, we apply the total variation loss on deformation feature grid to smooth the voxel features as

$$\mathcal{L}^{\text{d_tv}} = \frac{1}{|\mathcal{V}|} \sum_{\mathbf{v} \in \mathcal{V} c \in [N_c]} \sqrt{\Delta_x^2(\mathbf{v}, c) + \Delta_y^2(\mathbf{v}, c) + \Delta_z^2(\mathbf{v}, c)}, \quad (9)$$

where $\Delta_x^2(\mathbf{v}, c)$ represents the square difference between c -th element in voxel $\mathcal{V}(i, j, k)$ and voxel $\mathcal{V}(i+1, j, k)$, and the same for $\Delta_y^2(\mathbf{v}, c)$ and $\Delta_z^2(\mathbf{v}, c)$.

The overall loss function for can be written as

$$\mathcal{L} = \mathcal{L}^{\text{photo}} + w^{\text{ptc}} \cdot \mathcal{L}^{\text{ptc}} + w^{\text{bg}} \cdot \mathcal{L}^{\text{bg}} + w^{\text{d_norm}} \cdot \mathcal{L}^{\text{d_norm}} + w^{\text{d_tv}} \cdot \mathcal{L}^{\text{d_tv}}, \quad (10)$$

where w^{ptc} , w^{bg} , $w^{\text{df_norm}}$, and $w^{\text{df_tv}}$ are weights to balance each components in the final coarse loss.

Following [48], we use the same strategy for voxel allocation, model point sampling, and low-density initialization. Considering that the deformation is more sensitive to resolution compared with radiance (density and color), we set a higher grid resolution to the deformation module than the canonical module.

Progressive Training Typically, when a time step goes further away from the canonical time, the deformation between this time step and the canonical time step get larger. To reduce the learning difficulty, we use a progressive training strategy. The optimization starts with training images close to canonical time, and progressively adds images with further time steps.

Coarse-to-fine Optimization To speed up the training while maintaining the model capacity, we adopt a coarse-to-fine optimization procedure following [31, 48]. We first optimize a coarse model to roughly recover the deformation and the canonical space geometry. Then, we use the coarse model to locate the object region and filter out a large portion of empty space, after which a fine model is optimized to recover a more accurate and detailed deformation and geometry. In the coarse model, we do not model the view-dependent effect, and the feature dimension of color grid K' is set to 3, which directly corresponds to the RGB color. Based on the optimized coarse module, we apply empty space filtering strategies, including finding fine-stage bounding box and empty point filtering to speed up training. Also, we initialize the fine-stage model with model weight trained during coarse stage. We present details of our empty space filtering strategies and fine model design in the supplementary material.

5 Experiments

5.1 Dataset and Metrics

We evaluate our method on the D-NeRF dataset [39], which contains eight dynamic scenes with 360° viewpoint settings. Beside synthetic dataset, we also

Table 1. **Quantitative comparison.** We report LPIPS (lower is better) and PSNR/SSIM (higher is better) on eight dynamic scenes of the D-NeRF dataset.

Methods	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
D-NeRF (half ¹) [39]	25.03	0.951	0.069	31.29	0.974	0.027	29.26	0.965	0.117	38.93	0.990	0.103
NDVG (half ¹)	25.53	0.949	0.073	35.53	0.988	0.014	29.80	0.965	0.037	34.58	0.972	0.114
NDVG (w/o occ ²)	25.16	0.956	0.067	34.14	0.980	0.026	29.88	0.963	0.047	37.14	0.986	0.080
NDVG (w/o grid ³)	26.45	0.959	0.065	34.42	0.980	0.025	29.08	0.956	0.050	37.78	0.988	0.063
NDVG (w/o refine ⁴)	24.30	0.946	0.092	28.59	0.940	0.070	26.85	0.935	0.081	28.17	0.954	0.178
NDVG (w/o filter ⁵)	19.55	0.927	0.104	31.75	0.961	0.051	27.71	0.948	0.068	35.47	0.988	0.054
NDVG (full ⁶)	26.49	0.960	0.067	34.41	0.980	0.027	30.00	0.963	0.046	37.52	0.987	0.075

Methods	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
D-NeRF (half) [39]	21.64	0.839	0.165	31.76	0.977	0.040	32.80	0.982	0.021	32.80	0.981	0.037
NDVG (half)	25.23	0.931	0.049	30.15	0.967	0.047	34.05	0.983	0.022	29.45	0.960	0.078
NDVG (w/o occ)	24.77	0.935	0.059	32.57	0.979	0.031	18.78	0.899	0.112	30.87	0.973	0.044
NDVG (w/o grid)	24.18	0.916	0.078	31.64	0.976	0.034	32.99	0.980	0.027	30.64	0.971	0.044
NDVG (w/o refine)	23.30	0.851	0.167	27.35	0.940	0.079	29.80	0.965	0.051	26.13	0.920	0.150
NDVG (w/o filter)	22.75	0.887	0.140	28.58	0.952	0.067	32.36	0.976	0.035	28.19	0.957	0.077
NDVG (full)	25.04	0.940	0.053	32.62	0.978	0.033	33.22	0.979	0.030	31.25	0.974	0.040

¹ half: using half resolution of the original dataset images² w/o occ: not using occlusion reasoning³ w/o grid: not using deformation feature grid, only deform MLP⁴ w/o refine: only using coarse training stage⁵ w/o filter: not using coarse training, direct optimize fine module⁶ full: using full resolution of the original dataset images

conduct experiments on real scenes, proposed by HyperNeRF [36]. This dataset captures images of real dynamic scenes with a multi-view camera rig consisting of 2 phones with around a 16cm baseline. The dynamic motion consists of both rigid and non-rigid deformation with unbounded scenes. We use several metrics for the evaluation: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) [67].

5.2 Implementation Details

We set the expected voxel number to 1,664k and 190³ for the grid in Deformation Module in the coarse and fine stages, respectively. For Canonical Module, we set the expected voxel number to 1,024k and 160³. The light-weight MLP in the deformation module has 4 layers, each with a width of 64. The light-weight MLP in the canonical module has 3 layers, each with a width of 128.

When training with full resolution images, we train 10k and 20k iterations for the coarse and fine stages for all scenes. When training with half-resolution images, we reduce the iteration to 5k and 10k for coarse and fine stages. In terms of positional encoding, we set the frequency to 5 for position and time, and 4 for direction. We use the Adam optimizer [20] and sample 8,192 rays per iteration. More details of settings can be found in our supplementary material.

5.3 Comparisons

Quantitative evaluation on the dataset We first quantitatively compare the results in Table 1. To compare with D-NeRF[39], we set the same 400×400 image resolution, and present average results on each scene assessed by metrics which are mentioned above. According to Table 1, we could see that our method

Table 2. Quantitative comparison on real scenes.

Methods	Time	3D Printer		Broom		Chicken		Peel Banana		Mean	
		PSNR \uparrow	MS-SSIM \uparrow	PSNR \uparrow	MS-SSIM \uparrow	PSNR \uparrow	MS-SSIM \uparrow	PSNR \uparrow	MS-SSIM \uparrow	PSNR \uparrow	MS-SSIM \uparrow
NeRF [31]	~ hours	20.7	0.780	19.9	0.653	19.9	0.777	20.0	0.769	20.1	0.745
NV [27]	~ hours	16.2	0.665	17.7	0.623	17.6	0.615	15.9	0.380	16.9	0.571
NSFF [22]	~ hours	27.7	0.947	26.1	0.871	26.9	0.944	24.6	0.902	26.3	0.916
Nerfies [35]	~ hours	20.6	0.830	19.2	0.567	26.7	0.943	22.4	0.872	22.2	0.803
HyperNeRF [36]	~ hours	20.0	0.821	19.3	0.591	26.9	0.948	23.3	0.896	22.4	0.814
TiNeuVox [9]	30 mins	22.8	0.841	21.5	0.686	28.3	0.947	24.4	0.873	24.3	0.837
NDVG (Ours)	35 mins	22.4	0.839	21.5	0.703	27.1	0.939	22.8	0.828	23.3	0.823

Table 3. **Training time and rendering speed comparison.** We report these using the public code of D-NeRF [39] on the same device (RTX 3090 GPU) with our method. We include the mean PSNR across eight scenes in the D-NeRF dataset [39] for comparison of synthesized image quality. Our method could achieve good PSNR, while spend much less optimization time and have faster rendering speed.

Methods	PSNR \uparrow	Training Time (s/scene) \downarrow	Rendering Speed (s/img) \downarrow
NeRF(half) [†] [31]	19.00	60185	4.5
D-NeRF(half) [39]	30.02	99034	8.7
NDVG(half)	30.32	1380	0.4
NDVG(w/o refine)	26.73	708	2.6
NDVG(w/o filter)	27.85	2487	3.5
NDVG(full)	31.08	2087	1.7

[†] We use implementation of D-NeRF[39] to train NeRF on dynamic dataset

NDVG achieves comparable results with D-NeRF [39] for all three metrics. For real scenes, we test on four scenes, namely the *Peel Banana*, *Chicken*, *Broom* and *3D Printer* following HyperNeRF [36]. We follow the same settings of the experiments of TiNeuVox [9] and report the metrics of PSNR and MS-SSIM in Table 2 (results of other methods are taken from the TiNeuVox paper). As shown in Table 2, our method could achieve comparable or even better results and at least 10x faster, which clearly demonstrates the effectiveness of our method.

Training time and rendering speed comparison The key contribution of our work is to accelerate the optimization speed of novel view synthesis models on dynamic scenes. In Table 3, our method with the same half resolution setting with D-NeRF[39], achieves 70 \times faster convergence with an even higher average PSNR. Though our main purpose is to speed up training, the proposed method also has a reasonably fast rendering speed, compared to complete neural network based model. According to Table 3, our method has a 20 \times faster rendering speed compared with D-NeRF[39].

For real scenes in Table 2, compared with previous methods without acceleration which takes hours or even days to train, our method could finish training in 35 minutes which is at least 10x faster. Compared with the concurrent research TiNeuVox [9], which also aims to speed up training, we could achieve comparable results with the same training time, without using cuda acceleration for ray points sampling and total variation computation.

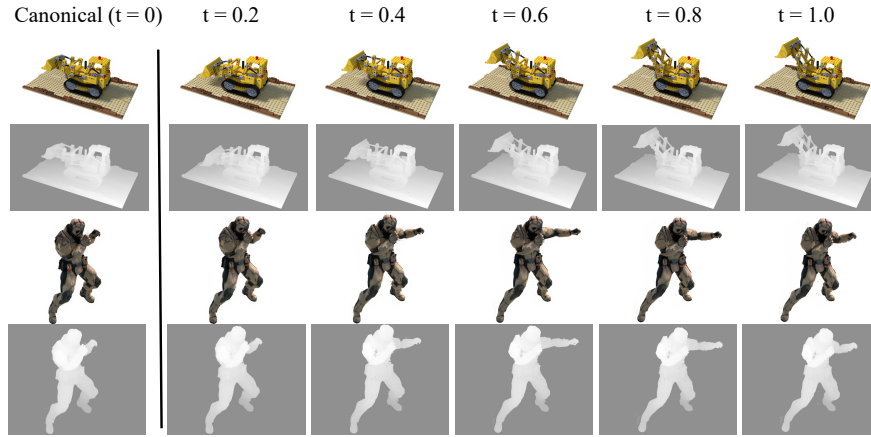


Fig. 4. **Learned Geometry.** We show examples of geometries learned by our model. For each, we show rendered images and corresponding disparity under two novel views and six time steps.

Qualitative comparison We provide some visualization of the learned scene representation in Fig. 4. We can see that our method can successfully recover the canonical geometry, and render high-quality dynamic sequences. The results indicate that our method can faithfully model the motion of the dynamic scenes.

In Fig. 5, we show the rendering results of more difficult situations and compare them with the results of D-NeRF [39]. Our method achieves comparable or even better image results using only 1/70 of the training time. If zoom in for more details, we could see that our method actually recovers more high-frequency details, taking the armour and the cloth of the worker as examples.

5.4 Method Analysis

In this section, we aim to study and prove the effectiveness of three designs in our proposed method: the occlusion reasoning, the coarse-to-fine optimization strategy, and the deformation feature grid.

Occlusion Reasoning We show that occlusion reasoning is critical in Sec. 3.4 for the canonical-based pipeline under the assumption that the empty point is static. Results of NDVG (full) and NDVG (w/o occ) in Table 1 compare the quantitative results of models with and without occlusion reasoning. We can see that the model with occlusion reasoning achieves more accurate results, verifying the effectiveness of our method.

We also visualize the estimated occlusion at different time steps in Fig. 6. We warp the canonical grid (density grid and color grid) into the different time steps, and show the points with corresponding colors if they are object points. And we show the occluded points estimated by deformation module in blue. In Fig. 6, while the character pushes out the punch and the body moves forward, the




Fig. 5. **Qualitative Comparison.** Synthesized images on test set of the dataset. For each scene, we show an image rendered at novel view, and followed by zoom in of ground truth, our NDVG, and D-NeRF [39].

empty points behind his back are estimated as being occluded which fits the actual situation. These results indicate the estimated occlusion is accurate.

Effectiveness of Coarse-to-fine Optimization The coarse-to-fine optimization could not only improve the image quality rendered by our model, but also help speed up the training and rendering process significantly. To prove the efficiency of the coarse-to-fine optimization, we conduct two extra experiments (see Table 1). The first one is NDVG (w/o refine), which only contains the coarse training stage. As NDVG (w/o refine) has limited grid resolution and does not have view-independent color representation, the performance is expected to be low. The second one is NDVG (w/o filter), which is not initialized by coarse training and directly begins from scratch for fine stage. As NDVG (w/o filter) does not have a trained coarse model for object region location, the fine model could not shrink the bounding box of the scene which means most of the grid space is wasted, which lead to obvious worse results. Also, the empty points cannot be filtered, which could increase training time and decrease rendering speed significantly, which is evidenced in Table 3.

Deformation Feature Grid In our deformation module, we use a deformation feature grid to encode dynamic information of 3D points, which will be decoded by a light-weight MLP to regress the deformation. A light-weight MLP is sufficient to model the complex motion of scenes by integrating with this feature grid. Table 1 shows that the method integrated with a deformation feature grid, NDVG (full), achieves better novel-view synthesis results than the one without, NDVG (w/o grid). This result clearly demonstrates the effectiveness of the proposed deformation feature grid. Please refer to our supplementary material for further study of the deformation feature grid design.

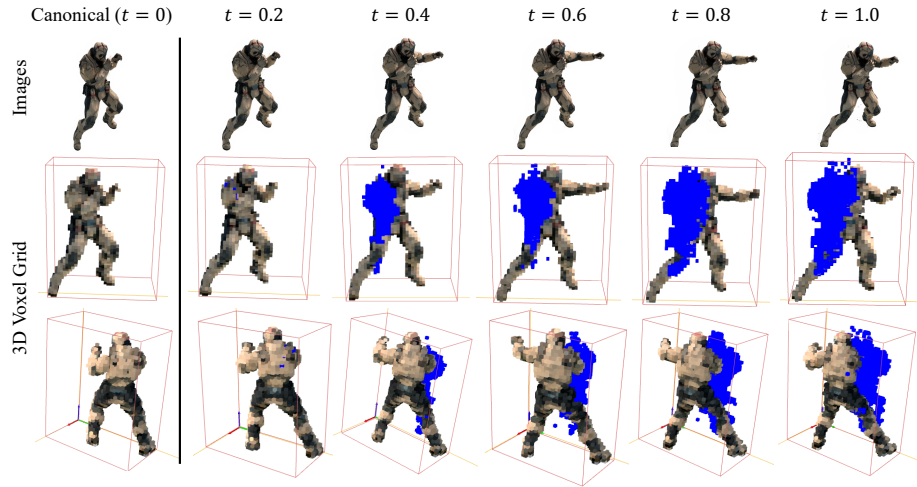


Fig. 6. **Occlusion Estimation Visualization.** We visualize the estimated occlusion points at different times in blue color. The first row shows the images at each time step which gives an insight of the motion. We warp the canonical grids into corresponding time step by the deformation estimated by deformation module. We visualize points with their rgb colors whose density is over 0.8.

6 Conclusions

In this paper, we have presented a fast optimization method for dynamic view synthesis based on a hybrid implicit-explicit representation. Our method consists of a deformation module to map a 3D point in observation space to canonical space, and a canonical module to represent the scene geometry and appearance. In each module, explicit dense grids and the light-weight MLP are used for fast feature interpolation and decoding, which significantly accelerates the optimization time compared with methods that rely on heavy MLP queries. Moreover, occlusion is explicitly modeled to improve the rendering quality. Experiments show that our method only requires 30 minutes to converge on a dynamic scene. Compared with the existing D-NeRF method, our method achieves a $70\times$ acceleration with comparable rendering quality.

Limitation Despite our method greatly speeds up the training of radiance field methods for dynamic view synthesis, it is mainly designed for bounded scenes. In the future, we will extend our method to deal with real-world unbounded scenes (*e.g.*, unbounded 360 and face-forward scenes).

Acknowledgements This work was supported in part by the National Natural Science Foundation of China (Nos. 61871325, 61901387, 62001394, 62202409) and the National Key Research and Development Program of China (No. 2018AAA0102803).

References

1. Aliev, K.A., Sevastopolsky, A., Kolos, M., Ulyanov, D., Lempitsky, V.: Neural point-based graphics. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 696–712 (2020) [3](#)
2. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 5855–5864 (2021) [3](#)
3. Boss, M., Braun, R., Jampani, V., Barron, J.T., Liu, C., Lensch, H.: Nerd: Neural reflectance decomposition from image collections. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 12684–12694 (2021) [3](#)
4. Buehler, C., Bosse, M., McMillan, L., Gortler, S., Cohen, M.: Unstructured lumigraph rendering. In: *ACM SIGGRAPH Computer Graphics (SIGGRAPH)*. pp. 425–432 (2001) [3](#)
5. Chen, A., Xu, Z., Zhao, F., Zhang, X., Xiang, F., Yu, J., Su, H.: MVSNeRF: Fast generalizable radiance field reconstruction from multi-view stereo. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 14124–14133 (2021) [2](#), [4](#)
6. Chen, S.E., Williams, L.: View interpolation for image synthesis. In: *ACM SIGGRAPH Computer Graphics (SIGGRAPH)*. pp. 279–288 (1993) [3](#)
7. Choi, I., Gallo, O., Troccoli, A., Kim, M.H., Kautz, J.: Extreme view synthesis. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 7781–7790 (2019) [3](#)
8. Du, Y., Zhang, Y., Yu, H.X., Tenenbaum, J.B., Wu, J.: Neural radiance flow for 4d view synthesis and video processing. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 14304–14314 (2021) [2](#), [4](#)
9. Fang, J., Yi, T., Wang, X., Xie, L., Zhang, X., Liu, W., Nießner, M., Tian, Q.: Fast dynamic radiance fields with time-aware neural voxels. *arXiv preprint arXiv:2205.15285* (2022) [4](#), [5](#), [11](#)
10. Flynn, J., Broxton, M., Debevec, P., DuVall, M., Fyffe, G., Overbeck, R., Snavely, N., Tucker, R.: Deepview: View synthesis with learned gradient descent. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2367–2376 (2019) [3](#)
11. Flynn, J., Neulander, I., Philbin, J., Snavely, N.: Deepstereo: Learning to predict new views from the world’s imagery. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5515–5524 (2016) [3](#)
12. Gan, W., Xu, H., Huang, Y., Chen, S., Yokoya, N.: V4d: Voxel for 4d novel view synthesis. *arXiv preprint arXiv:2205.14332* (2022) [4](#), [5](#)
13. Gao, C., Saraf, A., Kopf, J., Huang, J.B.: Dynamic view synthesis from dynamic monocular video. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 5712–5721 (2021) [2](#), [4](#)
14. Greene, N.: Environment mapping and other applications of world projections. *IEEE Computer Graphics and Applications (CG&A)* **6**(11), 21–29 (1986) [3](#)
15. Habtegebrial, T., Jampani, V., Gallo, O., Stricker, D.: Generative view synthesis: From single-view semantics to novel-view images. *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)* (2020) [3](#)
16. Hedman, P., Philip, J., Price, T., Frahm, J.M., Drettakis, G., Brostow, G.: Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (TOG)* **37**(6), 1–15 (2018) [3](#)

17. Huang, H.P., Tseng, H.Y., Lee, H.Y., Huang, J.B.: Semantic view synthesis. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 592–608 (2020) [3](#)
18. Kajiya, J.T., Von Herzen, B.P.: Ray tracing volume densities. ACM SIGGRAPH Computer Graphics (SIGGRAPH) **18**(3), 165–174 (1984) [2](#)
19. Kalantari, N.K., Wang, T.C., Ramamoorthi, R.: Learning-based view synthesis for light field cameras. ACM Transactions on Graphics (TOG) **35**(6), 1–10 (2016) [3](#)
20. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proceedings of the The International Conference on Learning Representations (ICLR) (2015) [10](#)
21. Levoy, M., Hanrahan, P.: Light field rendering. In: ACM SIGGRAPH Computer Graphics (SIGGRAPH). pp. 31–42 (1996) [3](#)
22. Li, Z., Niklaus, S., Snavely, N., Wang, O.: 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). pp. 6498–6508 (2021) [2](#), [4](#), [11](#)
23. Lin, C.H., Ma, W.C., Torralba, A., Lucey, S.: Barf: Bundle-adjusting neural radiance fields. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 5741–5751 (2021) [3](#)
24. Lindell, D.B., Martel, J.N., Wetzstein, G.: Autoint: Automatic integration for fast neural volume rendering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14556–14565 (2021) [4](#)
25. Liu, J.W., Cao, Y.P., Mao, W., Zhang, W., Zhang, D.J., Keppo, J., Shan, Y., Qie, X., Shou, M.Z.: Devrf: Fast deformable voxel radiance fields for dynamic scenes. arXiv preprint arXiv:2205.15723 (2022) [4](#), [5](#)
26. Liu, L., Gu, J., Lin, K.Z., Chua, T.S., Theobalt, C.: Neural sparse voxel fields. Proceedings of the Advances in Neural Information Processing Systems (NeurIPS) **33**, 15651–15663 (2020) [4](#)
27. Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehrmann, A., Sheikh, Y.: Neural volumes: Learning dynamic renderable volumes from images. ACM Transactions on Graphics (TOG) **38**(4), 65:1–65:14 (Jul 2019) [11](#)
28. Lombardi, S., Simon, T., Schwartz, G., Zollhoefer, M., Sheikh, Y., Saragih, J.: Mixture of volumetric primitives for efficient neural rendering. ACM Transactions on Graphics (TOG) **40**(4), 1–13 (2021) [4](#)
29. Martin-Brualla, R., Radwan, N., Sajjadi, M.S., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7210–7219 (2021) [3](#)
30. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. ACM Transactions on Graphics (TOG) **38**(4), 1–14 (2019) [3](#)
31. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 405–421 (2020) [1](#), [2](#), [3](#), [6](#), [9](#), [11](#)
32. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. arXiv preprint arXiv:2201.05989 (2022) [2](#), [4](#)
33. Neff, T., Stadlbauer, P., Parger, M., Kurz, A., Mueller, J.H., Chaitanya, C.R.A., Kaplanyan, A.S., Steinberger, M.: DONeRF: Towards Real-Time Rendering of

- Compact Neural Radiance Fields using Depth Oracle Networks. *Computer Graphics Forum* **40**(4) (2021) [4](#)
34. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3504–3515 (2020) [1](#)
 35. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (2021) [2](#), [4](#), [11](#)
 36. Park, K., Sinha, U., Hedman, P., Barron, J.T., Bouaziz, S., Goldman, D.B., Martin-Brualla, R., Seitz, S.M.: HyperNeRF: A higher-dimensional representation for topologically varying neural radiance fields. *ACM Transactions on Graphics (TOG)* **40**(6), 1–12 (2021) [5](#), [10](#), [11](#)
 37. Penner, E., Zhang, L.: Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)* **36**(6), 1–11 (2017) [3](#)
 38. Píala, M., Clark, R.: Terminerf: Ray termination prediction for efficient neural rendering. In: *Proceedings of the International Conference on 3D Vision (3DV)*. pp. 1106–1114 (2021) [4](#)
 39. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 10318–10327 (2021) [2](#), [3](#), [4](#), [5](#), [9](#), [10](#), [11](#), [12](#), [13](#)
 40. Rebain, D., Jiang, W., Yazdani, S., Li, K., Yi, K.M., Tagliasacchi, A.: Derf: Decomposed radiance fields. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 14153–14161 (2021) [4](#)
 41. Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. pp. 14335–14345 (2021) [4](#)
 42. Riegler, G., Koltun, V.: Free view synthesis. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 623–640 (2020) [3](#)
 43. Riegler, G., Koltun, V.: Stable view synthesis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 12216–12225 (2021) [3](#)
 44. Shum, H., Kang, S.B.: Review of image-based rendering techniques. In: *Visual Communications and Image Processing (VCIP)*. vol. 4067, pp. 2–13 (2000) [3](#)
 45. Srinivasan, P.P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B., Barron, J.T.: Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7495–7504 (2021) [3](#)
 46. Srinivasan, P.P., Mildenhall, B., Tancik, M., Barron, J.T., Tucker, R., Snavely, N.: Lighthouse: Predicting lighting volumes for spatially-coherent illumination. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8080–8089 (2020) [3](#)
 47. Srinivasan, P.P., Tucker, R., Barron, J.T., Ramamoorthi, R., Ng, R., Snavely, N.: Pushing the boundaries of view extrapolation with multiplane images. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 175–184 (2019) [3](#)
 48. Sun, C., Sun, M., Chen, H.T.: Direct Voxel Grid Optimization: Super-fast convergence for radiance fields reconstruction. *arXiv preprint arXiv:2111.11215* (2021) [2](#), [4](#), [5](#), [7](#), [8](#), [9](#)

49. Tewari, A., Fried, O., Thies, J., Sitzmann, V., Lombardi, S., Xu, Z., Simon, T., Nießner, M., Tretschk, E., Liu, L., et al.: Advances in neural rendering. In: ACM SIGGRAPH Computer Graphics (SIGGRAPH). pp. 1–320 (2021) [3](#)
50. Tewari, A., Fried, O., Thies, J., Sitzmann, V., Lombardi, S., Sunkavalli, K., Martin-Brualla, R., Simon, T., Saragih, J., Nießner, M., et al.: State of the art on neural rendering. In: Computer Graphics Forum. pp. 701–727 (2020) [3](#)
51. Thies, J., Zollhöfer, M., Nießner, M.: Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)* **38**(4), 1–12 (2019) [3](#)
52. Tretschk, E., Tewari, A., Golyanik, V., Zollhöfer, M., Lassner, C., Theobalt, C.: 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). pp. 12959–12970 (2021) [2](#), [4](#), [5](#)
53. Trevithick, A., Yang, B.: Grf: Learning a general radiance field for 3d representation and rendering. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 15182–15192 (2021) [2](#), [4](#)
54. Tucker, R., Snaveley, N.: Single-view view synthesis with multiplane images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 551–560 (2020) [3](#)
55. Wang, C., Eckart, B., Lucey, S., Gallo, O.: Neural trajectory fields for dynamic novel view synthesis. *arXiv preprint arXiv:2105.05994* (2021) [2](#), [4](#)
56. Wang, Q., Wang, Z., Genova, K., Srinivasan, P.P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snaveley, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4690–4699 (2021) [2](#), [4](#)
57. Wang, Z., Wu, S., Xie, W., Chen, M., Prisacariu, V.A.: Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064* (2021) [3](#)
58. Xian, W., Huang, J.B., Kopf, J., Kim, C.: Space-time neural irradiance fields for free-viewpoint video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9421–9431 (2021) [2](#), [4](#)
59. Xu, Z., Bi, S., Sunkavalli, K., Hadap, S., Su, H., Ramamoorthi, R.: Deep view synthesis from sparse photometric images. *ACM Transactions on Graphics (TOG)* **38**(4), 1–13 (2019) [3](#)
60. Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Basri, R., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. In: Proceedings of the Advances in Neural Information Processing Systems (NeurIPS). vol. 33, pp. 2492–2502 (2020) [1](#)
61. Yen-Chen, L., Florence, P., Barron, J.T., Rodriguez, A., Isola, P., Lin, T.Y.: inerf: Inverting neural radiance fields for pose estimation. In: Proceedings of the International Conference on Intelligent Robots and Systems (IROS). pp. 1323–1330 (2021) [3](#)
62. Yoon, J.S., Kim, K., Gallo, O., Park, H.S., Kautz, J.: 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). pp. 5336–5345 (2020) [4](#)
63. Yu, A., Fridovich-Keil, S., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#), [5](#)

64. Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenotrees for real-time rendering of neural radiance fields. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 5752–5761 (2021) [2](#), [4](#)
65. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4578–4587 (2021) [2](#), [4](#)
66. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: NERF++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492 (2020) [3](#)
67. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 586–595 (2018) [10](#)
68. Zhang, X., Srinivasan, P.P., Deng, B., Debevec, P., Freeman, W.T., Barron, J.T.: Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. ACM Transactions on Graphics (TOG) **40**(6), 1–18 (2021) [3](#)
69. Zhou, T., Tucker, R., Flynn, J., Fyfe, G., Snavely, N.: Stereo magnification: learning view synthesis using multiplane images. ACM Transactions on Graphics (TOG) **37**(4), 1–12 (2018) [3](#)