

Fourier PlenOctrees for Dynamic Radiance Field Rendering in Real-time

Liao Wang^{1*} Jiakai Zhang^{1*} Xinhang Liu¹ Fuqiang Zhao¹ Yanshun Zhang³
 Yingliang Zhang³ Minye Wu² Jingyi Yu^{1#} Lan Xu^{1,4#}
¹ ShanghaiTech University ² KU leuven ³ DGene
⁴ Shanghai Engineering Research Center of Intelligent Vision and Imaging

Abstract

Implicit neural representations such as Neural Radiance Field (NeRF) have focused mainly on modeling static objects captured under multi-view settings where real-time rendering can be achieved with smart data structures, e.g., PlenOctree. In this paper, we present a novel Fourier PlenOctree (FPO) technique to tackle efficient neural modeling and real-time rendering of dynamic scenes captured under the free-view video (FVV) setting. The key idea in our FPO is a novel combination of generalized NeRF, PlenOctree representation, volumetric fusion and Fourier transform. To accelerate FPO construction, we present a novel coarse-to-fine fusion scheme that leverages the generalizable NeRF technique to generate the tree via spatial blending. To tackle dynamic scenes, we tailor the implicit network to model the Fourier coefficients of time-varying density and color attributes. Finally, we construct the FPO and train the Fourier coefficients directly on the leaves of a union PlenOctree structure of the dynamic sequence. We show that the resulting FPO enables compact memory overload to handle dynamic objects and supports efficient fine-tuning. Extensive experiments show that the proposed method is 3000 times faster than the original NeRF and achieves over an order of magnitude acceleration over SOTA while preserving high visual quality for the free-viewpoint rendering of unseen dynamic scenes.

1. Introduction

Interactive and immersive applications, such as Telepresence and Virtual Reality (VR), make plenty use of free-viewpoint videos to provide unique and fully controllable viewing experiences. At the core are fast generation and real-time rendering at new viewpoints with ultra-high pho-

* Authors contributed equally to this work. # The corresponding authors are Jingyi Yu (yujingyi@shanghaitech.edu.cn) and Lan Xu (xulan1@shanghaitech.edu.cn).



Figure 1. Our method achieves a photo-realistic rendered result for dynamic scenes in real-time based on our novel Fourier PlenOctree structure.

to-realism. Traditional image-based modeling and rendering approaches rely on feature matching and view interpolation, whereas the latest neural rendering techniques are able to integrate the two processes into a deep net that simultaneously represents the geometry and appearance for efficient rendering. By far, most neural rendering techniques have focused on modeling static objects and employing smart data structures. For example, volumetric neural modeling techniques [28, 32] overcome many limitations of traditional methods, including tailored matching algorithms and optimization procedures and can even tackle non-Lambertian materials. The seminal work of the Neural Radiance Field (NeRF) [32] learns a neural representation based on MLP to represent static scenes as radiance fields with the property of density and color. It only requires calibrated multi-view images to produce compelling free-viewpoint rendering. However, the MLP structure is still too slow to achieve real-time performance. Existing techniques explore using thousands of tiny MLPs [45], applying factorization [13], tailored volumetric data structures [16, 69], and primitive-based rendering [29]. Despite their effectiveness, very few techniques are directly applicable to handle dynamic scenes, in particular, objects with non-rigid deformations such as the human body. In this work, we present a novel neural representation for generating free-viewpoint videos from multi-view

sequence inputs as well as for real-time photorealistic rendering.

Several recent efforts have investigated combining neural modeling with classical flow or geometry estimations [38, 43]. For example, one can apply motion flows to explicitly transform sampled points in individual frames to a canonical model to partially account for non-rigid deformations. However, they are vulnerable to lost tracks and can lead to motion artifacts. There are also emerging interests on utilizing pre-defined (e.g. skeleton [40] or parametric models [27, 41]) to explicitly calculate stable motion flows from model animations. These approaches, however, are limited to handling specific types of objects consistent with the pre-defined model. [58] directly predicts a neural radiance field using a general network for each frame while avoiding online training. Its rendering speed, however, is not yet sufficient for interactive and immersive experiences.

In this paper, we present a novel Fourier PlenOctree (FPO) technique for neural dynamic scene representation, which enables efficient neural modeling and real-time rendering of unseen dynamic objects with compact memory overload, as illustrated in Fig. 1. The key idea in our FPO is a novel combination of generalized NeRF, PlenOctree representation, volumetric fusion, and Fourier transform.

For efficient scene modeling, we present a novel coarse-to-fine fusion scheme that leverages generalizable NeRF [59] technique to generate the PlenOctree with fast plenoptic functions inference. Inspired by the volumetric fusion strategy [36], we propose a spatial blending scheme to generate the neural tree in the order of minutes. To tackle dynamic scenes, we tailor the implicit network to model the Fourier coefficients of time-varying density and plenoptic functions of the dynamic PlenOctree. By discarding high-frequency bases, our novel representation can achieve high storage efficiency while persevering perceptive details. Finally, we construct the FPO and train the Fourier coefficients directly on the leaves of a union PlenOctree structure of the dynamic sequence. By combining the benefits of PlenOctree rendering and Fourier operations, our FPO enables real-time free-viewpoint synthesis of dynamic scenes and supports efficient fine-tuning. Comprehensive experiments show that FPO is 3000 times faster than the original NeRF implementation and achieves over an order of magnitude acceleration over state-of-the-art techniques for dynamic scene modeling and rendering. To summarize, our contributions include:

- We introduce a FPO representation that enables real-time rendering of general dynamic scenes with fast fine-tuning and compact memory overload.
- We present a coarse-to-fine scheme that utilizes generalizable NeRF for PlenOctree generation and constructing FPO efficiently.

2. Related work

Novel View Synthesis for Static Scenes. The task of synthesizing novel views of a scene given a set of photographs has attracted the attention of the community. A great many current methods predict an underlying geometric or image-based 3D representation that allows rendering from novel viewpoints. To tackle novel view synthesis, [14, 22] formalize the concept of a light field. [48] represents both geometry and appearance of the underlying 3D scene in a 360-degree, four-dimensional light field parameterized via a neural implicit representation. Image-based representations [4, 67] denote any set of images conditioned on combinations of time, view, light or other dimensions like the color spectrum. Mesh-based representations [57, 61] are compact and easy to render; however, optimizing a mesh to fit a complex scene of arbitrary topology is challenging. Volume rendering is a classical technique with a long history of research in the graphics community [9]. Volume-based representations such as voxel grids [21, 47] and multi-plane images (MPIs) [42, 51] are a popular alternative to mesh representations due to their topology-free nature: gradient-based optimization is therefore straightforward, while rendering can still be real-time.

The most notable approach Neural Radiance Field (NeRF) [32] combines the implicit representation with volumetric rendering to achieve compelling novel view synthesis with rich view-dependent effects. However, these neural representations above can only handle static scenes.

Novel View Synthesis for Dynamic Scenes. Different from static scenes, tackling dynamic scenes requires settling the illumination changes and moving objects. One approach is to obtain a reconstruction for dynamic objects with input data from either camera array or a single camera. Methods involving silhouette [19, 53], stereo [23, 30, 31, 64], segmentation [44, 46], and photometric [1, 15, 56] have been explored. Early solutions [6, 7, 34] rely on multi-view dome-based setup for high-fidelity reconstruction and texture rendering of human activities in novel views. Recently, volumetric approaches with RGB-D sensors and modern GPUs have enabled real-time novel view synthesis for dynamic scenes and eliminated the reliance on a pre-scanned template model. The high-end solutions [7, 8, 18, 65] rely on multi-view studio setup to achieve high-fidelity reconstruction and rendering, while the low-end approaches [17, 35, 49, 63] adopt the most handy monocular setup with a temporal fusion pipeline [36] but suffer from inherent self-occlusion constraint.

Recent work [2, 10, 11, 24, 37–39, 43, 54, 62, 71] extend the approach NeRF [32] using neural radiance field into the dynamic settings. Some [10, 11, 62] condition the neural radiance field on time in addition to space, and others [24, 38, 43, 54, 71, 73] decompose the task into learning a spatial mapping from the current scene to the canonical

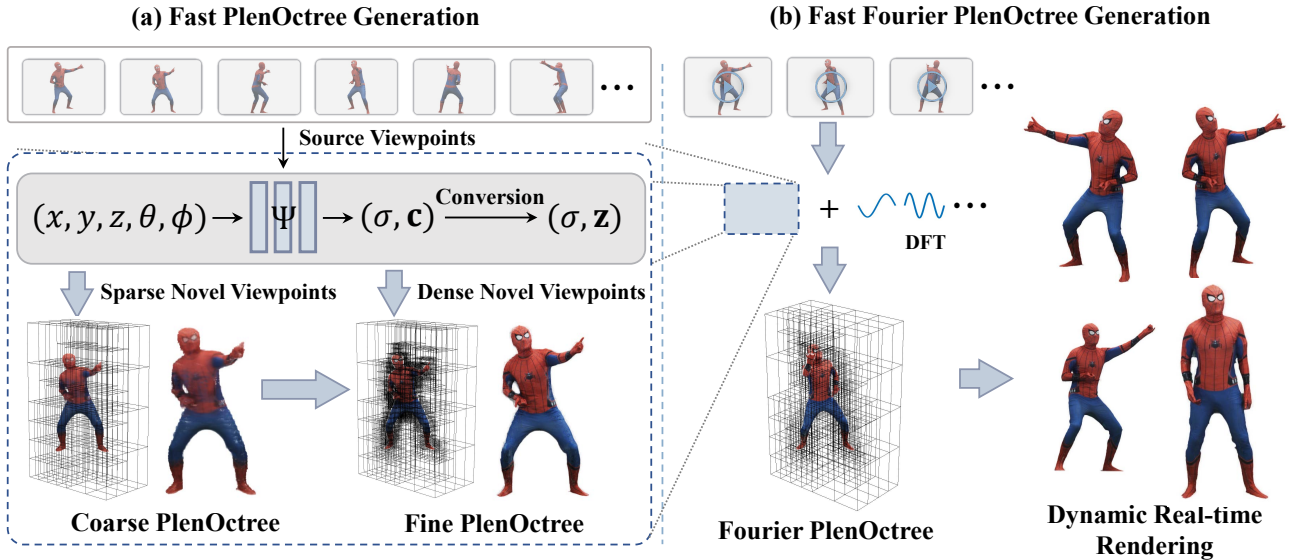


Figure 2. Illustration of our fast PlenOctree generation for static scene and fast Fourier PlenOctree generation for dynamic scenes. (a) illustrates how to generate a PlenOctree from multiview images. A generalized NeRF Ψ predicts view-dependent density σ and color \mathbf{c} by inputting 3D sample point (x, y, z) with view direction (θ, ϕ) and source viewpoints, then we can convert them to view-independent density σ and SH coefficients \mathbf{z} . Using sparse novel view RGB images rendered by Ψ , we can obtain a coarse PlenOctree. Finally, we fine-tune it to be a fine PlenOctree by using dense novel view images. (b) extends the pipeline to dynamic scenes by combining the idea with Discrete Fourier Transform (DFT) and achieves a fast Fourier PlenOctree generation and real-time rendering for dynamic scenes.

scene at each timestamp and regressing the canonical radiance field. [37] uses a learned scene graph with decoupled object transformations and scene representations to represent dynamic, multi-object scenes. [39] conditions NeRF on additional higher-dimensional coordinates, which can be thought of as a higher dimensional analog to the deformation field. [2] uses the depth information captured with time-of-flight (ToF) cameras in addition to RGB images to optimize for a neural radiance field. However, the above solutions using dynamic neural radiance fields still suffer from a long training time as well as rendering time.

NeRF Accelerations. While NeRFs can produce high-quality results, their computationally expensive rendering leads to slow training and inference. One way to speed up the process of fitting a NeRF to a new scene is to incorporate priors learned from a dataset of similar scenes. This can be accomplished by conditioning on predicted images features [55, 59, 70] or meta-learning [52]. To improve rendering speed, Neural Sparse Voxel Fields (NSVF) [26] learns sparse voxel grids of features that are input into a NeRF-like model. The sparse voxel grid allows the renderer to skip over empty regions when tracing a ray which improves the rendering time $\sim 10x$. AutoInt [25] modifies the architecture of the NeRF so that inference requires fewer samples but produces lower quality results.

NeX [60] extends MPIs to encode spherical basis functions that enable view-dependent rendering effects in real-time. [13, 16, 45] also distill NeRFs to enable real-time rendering. [69] use an octree-based 3D representation which

supports view-dependent effects to achieve real-time performance. Recent concurrent works [20, 33, 50, 66, 68] can accelerate radiance field training by an order of magnitude, but they can only handle static scenes.

However, none of the current methods tackles the challenge to accelerate the training and rendering process of the dynamic radiance field.

3. Generalized PlenOctree Fusion

Recall that NeRF takes an MLP as a mapping function to predict a density σ and a color \mathbf{c} for a queried 3D point $p = (x, y, z)$ in a given viewing direction $\mathbf{d} = (\theta, \phi)$. To boost the NeRF rendering procedure, [69] modifies the outputs of the mapping function to Spherical Harmonic (SH) coefficients $\mathbf{z} \in \mathbb{R}^{\ell_{max} \times 3}$ with a density σ , which will be cached in leaves of PlenOctree as an initialization. Having \mathbf{z} and σ , we can calculate the color of queried 3D point in given viewing direction:

$$\mathbf{c}(\mathbf{d}; \mathbf{z}) = S \left(\sum_{\ell=0}^{\ell_{max}} \sum_{m=-\ell}^{\ell} z_{m,\ell} Y_{m,\ell}(\mathbf{d}) \right) \quad (1)$$

where S is the sigmoid function to normalize color, $Y_{m,\ell} : S^2 \rightarrow \mathbb{R}$ is a real-valued basis function of SH.

Even though the rendering speed of PlenOctree is rather fast due to this simple calculation, the acquisition of SH coefficients and densities is still time-consuming. Therefore, we present a novel coarse-to-fine fusion scheme that leverages the generalizable NeRF technique Ψ [5, 58, 59]

to attack this problem via spatial blending. In the following, we demonstrate PlenOctree Fusion algorithm in traditional static scenes as an example. Note that we can also do PlenOctree Fusion in Fourier PlenOctree to deal with dynamic scenarios in the same way since both of them have very similar data structures which will be introduced in Sec. 4.2.

A generalized neural rendering network Ψ takes images of adjacent views of a target view as inputs and infers an implicit volume with regard to the target view. We can directly query colors and densities of sample points corresponding to leaves in PlenOctree from this volume without per-scene training. However, these colors and densities are all with respect to a specific target view due to different view directions. To obtain a completed PlenOctree, we need to sample more target views and fuse their local PlenOctree together. Fig. 2 (a) illustrates our pipeline.

The proposed PlenOctree Fusion follows a coarse-to-fine strategy. Specifically, to obtain coarse PlenOctrees, we initialize a coarse PlenOctree with N^3 voxel grids as tree leaves. Given multi-view images and silhouettes extracted via chroma key segmentation and background subtraction, Ψ predicts images for 6 sparse views which are uniformly around the object. Then Shape-From-Silhouette (SFS) method [3] is applied to generate a coarse visual hull from sparse view silhouettes. For each leaf inside the visual hull, we uniformly sample directions $\theta, \phi \in [0, 2\pi]$ to predict densities and colors by feeding positions and directions to Ψ . Note that predicted densities and colors are both view-dependent, denoted as $\sigma(\theta, \phi)$ and $\mathbf{c}(\theta, \phi)$ respectively. Next, we need to convert those predictions to view-independent densities and SH coefficients \mathbf{z} for each leaf by:

$$\begin{aligned} \sigma &= \mathbb{E}(\sigma(\theta, \phi)) \\ \mathbf{z} &= \mathbb{E}(\text{SH}(\theta, \phi) \cdot \mathbf{c}(\theta, \phi)), \end{aligned} \quad (2)$$

where $\mathbb{E}(\cdot)$ is the expectation operator, and $\text{SH}(\theta, \phi)$ is an evaluation function which calculates SH from a direction. After filling in all leaves, we obtain a coarse PlenOctree.

In the fine stage, inspired by fusion-based methods [12, 36], we first render 100 dense view images by PlenOctree, and query the points that the transmittance in volume rendering $T_i > 1e - 3$ from Ψ , then we fuse PlenOctree using these points by calculating the updating weights for each leaf. The reason why PlenOctree is initialized from 6 views is that the 6 views query all the points, while the 100 views will only query about 1% points which are fast compared to querying all the points of 100 views. At i -th viewpoints, we use the transmittance $T_i(x, y, z)$ as update weight for leaves at position (x, y, z) and update density and SH coefficients by the following equation:

$$\sigma_i = \frac{W_{i-1}\sigma_{i-1} + T_i\sigma_i}{W_{i-1} + T_i} \quad (3)$$

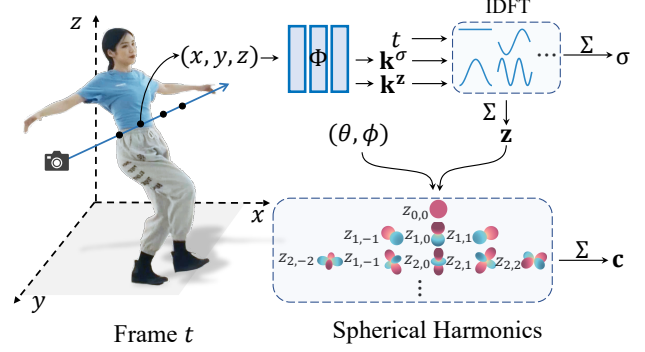


Figure 3. Illustration of our 4D Scene Representation in frequency domain, including first predicting Fourier coefficients \mathbf{k}^σ and \mathbf{k}^z by inputting (x, y, z) , then computing density σ and factors \mathbf{z} of SH basis by summing the weighted Fourier Transform with additional timestamp t , finally, computing color \mathbf{c} by summing the weighted SH bases with viewing direction (θ, ϕ) .

$$\mathbf{z} = \frac{W_{i-1}\mathbf{z}_{i-1} + T_i \cdot \text{SH}(\theta_i, \phi_i) \cdot \mathbf{c}(\theta_i, \phi_i)}{W_{i-1} + T_i} \quad (4)$$

Then calculate weight and count of updated times:

$$W_i = \frac{C_i - 1}{C_i} W_{i-1} + \frac{1}{C_i} T_i \quad (5)$$

where $C_i = C_{i-1} + 1$ means how many times the leaf has been updated at i -th step. After these iterations, we filter out the leaves which have $\sigma < 1e - 3$ to save the PlenOctree storage and further computational costs.

With the help of the generalized PlenOctree Fusion, we can obtain a PlenOctree representation for a static scene within 60 seconds, which greatly improves generation speed.

4. Fourier PlenOctree

In this section, we elaborate how Fourier PlenOctree records and renders free-viewpoint videos. As illustrated in Fig. 2 (b), we introduce Fourier PlenOctree with a novel 4D scene representation, which adopts PlenOctree to dynamic scenes by compressing time-variant information in the frequency domain (Sec. 4.1). Combined with Generalized PlenOctree Fusion, Fourier PlenOctree exhibits fast generation and real-time rendering abilities (Sec. 4.2). Fourier PlenOctree fine-tuning can further improve rendering quality within a short time (Sec. 4.3).

4.1. 4D Scene Representation in Frequency Domain

We propose a novel 4D Scene Representation in a high dimensional frequency domain to enable efficient neural modeling and real-time rendering for general dynamic scenes.

As illustrated in Fig. 3, given a 4D scene sample point (x, y, z, t) , the mapping function Φ is defined as below:

$$\Phi(x, y, z) = \mathbf{k}^\sigma, \mathbf{k}^z \quad (6)$$

where $\mathbf{k}^\sigma \in \mathbb{R}^{n_1}$ and $\mathbf{k}^z \in \mathbb{R}^{n_2 \times (\ell_{max} + 1)^2 \times 3}$ are two Fourier Transform coefficients of the functions of density $\sigma(t)$ and SH coefficients $\mathbf{z}(t)$ at position (x, y, z) respectively; n_1 is the Fourier coefficient number of σ , n_2 is the Fourier coefficient number of each SH coefficient \mathbf{z} , note that $(\ell_{max} + 1)^2 \times 3$ is the number of SH coefficients for RGB channels. As the timestamp t is given, density σ can be recovered by the following real-valued Inverse Discrete Fourier Transform (IDFT) in Eq. 7:

$$\sigma(t; \mathbf{k}^\sigma) = \sum_{i=0}^{n_1-1} \mathbf{k}_i^\sigma \cdot \text{IDFT}_i(t) \quad (7)$$

where t is the frame index and

$$\text{IDFT}_i(t) = \begin{cases} \cos\left(\frac{i\pi}{T}t\right) & \text{if } i \text{ is even} \\ \sin\left(\frac{(i+1)\pi}{T}t\right) & \text{if } i \text{ is odd} \end{cases} \quad (8)$$

To handle view-dependent effects, We use a similar idea to compute Fourier coefficients for each element $z_{m,l} \in \mathbb{R}^3$ of coefficients $\mathbf{z} = (z_{m,l})_{\substack{m: -\ell \leq m \leq \ell \\ l: 0 \leq l \leq \ell_{max}}}$ of SH function by the following Eq. 9:

$$z_{m,\ell}(t; \mathbf{k}^z) = \sum_{i=0}^{n_2-1} \mathbf{k}_{m,\ell,i}^z \cdot \text{IDFT}_i(t) \quad (9)$$

where $\mathbf{k}_{m,\ell,i}^z \in \mathbb{R}^3$ is defined as $\mathbf{k}^z = (k_{m,\ell,i})_{\substack{m: -\ell \leq m \leq \ell \\ l: 0 \leq l \leq \ell_{max}}}$ and $0 \leq i \leq n_2$ is an additional dimension to store corresponding Fourier coefficients for each SH element $z_{m,l}$.

Similar to PlenOctree [69], the mapping function Φ can be adapted by an MLP based NeRF network which we called Fourier NeRF-SH and be further discretized into octree-based volume representation. Content in each leaf contains the Fourier coefficients \mathbf{k}^σ and \mathbf{k}^z of the corresponding position. As a result, the proposed representation absorbs the advantages and benefits of the original PlenOctree and enables real-time novel view synthesis for free-viewpoint videos.

4.2. Fourier PlenOctree Generation

Reconstruction a Fourier PlenOctree as described in Sec. 4.1 is a big challenge. A naive way to reconstruct such Fourier PlenOctree is to fit a continual implicit function as described in Eq. 6 from scratch using Fourier NeRF-SH like [69], which takes about 1-2 days. For speed considerations, we adopt Generalized PlenOctree Fusion (Section. 3) in free-viewpoint video generation.

Octree structures vary from frame to frame due to object's motion. Fourier PlenOctree requires the structures to

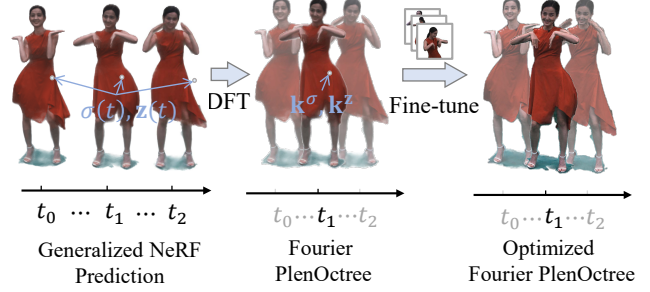


Figure 4. A straightforward pipeline to generate Fourier PlenOctree from multiple PlenOctrees, for each leaf in Fourier PlenOctree, we first find all corresponding PlenOctree leaves at the same position. They store a sequence of density σ and SH coefficients $\mathbf{z}(t)$ along the time axis. We can convert them to Fourier coefficients \mathbf{k}^σ and \mathbf{k}^z corresponding to density and SH coefficients, and store them in Fourier PlenOctree. Finally, we can optimize Fourier PlenOctree using ground truth images.

be the same in all frames in order to analyze plenoptic functions located at the same position. Fortunately, we are able to fast infer octree structures via Shape-From-Silhouette (SFS). Applying Generalized PlenOctree Fusion for each frame, we fill content in all frames' PlenOctrees. After that, all we need is to unify them. For PlenOctrees at timestamps $t = 1, 2, \dots, T$, we first calculate the union of their structures, note that the union of their structures always has equal or deeper leaves comparing PlenOctree at any frame. In other words, each leaf in an octree either is divided or keeps the same. In the case of division, we just copy the pre-existed value from the parent node (previous leaf) to new leaves.

Having unified PlenOctrees for each frame, we calculate a Fourier PlenOctree, as shown in Fig. 4, which has the same octree structure as theirs by using the Discrete Fourier Transform (DFT) for each leaf's values \mathbf{k}^σ and \mathbf{k}^z :

$$\mathbf{k}_i^\sigma = \sum_{t=1}^T \sigma(t) \cdot \text{DFT}_i(t) \quad (10)$$

$$\mathbf{k}_{m,\ell,i}^z = \sum_{t=1}^T z_{m,\ell}(t) \cdot \text{DFT}_i(t) \quad (11)$$

where

$$\text{DFT}_i(t) = \begin{cases} \frac{1}{T} \cos\left(\frac{i\pi}{T}t\right) & \text{if } i \text{ is even} \\ \frac{1}{T} \sin\left(\frac{(i+1)\pi}{T}t\right) & \text{if } i \text{ is odd} \end{cases} \quad (12)$$

4.3. Fourier PlenOctree Fine-tuning

Although our proposed Fourier PlenOctree has a DFT mapping function from Fourier coefficients to densities and SH coefficients at a specific timestamp, the fine-tuning procedure discussed in [69] can be extended to our method to



Figure 5. The rendered appearance results of our Fourier PlenOctree method on several sequences, including humans, human with objects and animals.

improve the image quality via back propagation as DFT is totally differentiable.

The objective function of this procedure is the same as the loss function of [32]:

$$\mathcal{L} = \sum_{t=1}^T \sum_{i=1}^N \|\hat{I}_i^t - I_i^t\|_2^2 \quad (13)$$

where \hat{I}_i^t is the rendered image for view i and frame t .

The optimization time is much shorter than optimizing Fourier NeRF-SH, since Fourier PlenOctree is an explicit representation which is easier to optimize than MLP-based implicit representations.

5. Experimental Results

In this section, we evaluate our Fourier PlenOctree method on a variety of challenging scenarios. We run our experiments on a PC with a single NVIDIA GeForce RTX3090 GPU. It only takes us about 2 hours to reconstruct Fourier PlenOctree with input from 60 views and 60 frames. For dynamic datasets, we have five real datasets in 2048×1536 and five synthetic datasets in 1080×1080 . The cameras are the cylindrical distribution looking at the object. We use $\ell_{max} = 2$ (9 components) and 512^3 grid size for our Fourier PlenOctree. Our method achieves rendering speed at 100fps with 800×800 resolution, which is 3000

times faster than the original NeRF. As demonstrated in Fig. 5, our approach generates high-quality appearance results and even handles identities with rich textures and challenging motions. Please refer to the supplementary video for more video results.

5.1. Comparison

Our approach is first to enable the fast generation and real-time dynamic rendering to the best of our knowledge. To demonstrate the overall performance of our approach, we compare to the existing free-viewpoint video methods based on neural rendering, including the voxel-based method **Neural Volumes** [28], and implicit methods **iButter** [58], **ST-NeRF** [71] and **Neural Body** [41] based on neural radiance field. For a fair comparison, all the methods share the same training dataset as our approach.

As shown in Fig. 6, our approach achieves photo-realistic free-viewpoint rendering with the most vivid rendering result in terms of photo-realism and sharpness, which, in addition, can be done in real-time.

For quantitative comparison, we adopt the peak signal-to-noise ratio (**PSNR**), structural similarity index (**SSIM**), mean absolute error (**MAE**), and Learned Perceptual Image Patch Similarity (**LPIPS**) [72] as metrics to evaluate our rendering accuracy. We keep 90 percent of captured views as training set and the other 10 percent views as testing set. As shown in Tab. 1, our approach outperforms




Figure 6. Qualitative comparison against dynamic scene rendering methods. We compare our method with Neural Body, Neural volumes and ST-NeRF and our approach generalizes the most photo-realistic and finer detail.

Method	best	second-best	MAE↓	LPIPS↓
	PSNR↑	SSIM↑		
Neural Body	27.34	0.9414	0.0123	0.0373
NeuralVolumes	23.62	0.9219	0.0251	0.0881
ST-NeRF	30.63	0.9486	0.0092	0.0324
iButter	33.77	0.9596	0.0054	0.0295
Ours	35.21	0.9910	0.0033	0.0217

Table 1. **Quantitative comparison against several methods in terms of rendering accuracy.** Compared with ST-NeRF, Neural Volumes, NeuralBody and iButter, our approach achieves the best performance in PSNR, SSIM, LPIPS and MAE metrics.

other methods in terms of all the metrics for appearance. Such a qualitative comparison illustrates the effectiveness of our approach to encode the spatial and temporal information from our multi-view setting. In Tab. 2, our method achieves the fastest rendering in dynamic scenes and uses the second least training or fine-tuning time given a new multi-view sequence.

5.2. Ablation Study

Time of fine-tuning. We analyze the relationship between rendering quality and fine-tuning time. In these experiments, we compared the results obtained without fine-tuning, with a 10-minute fine-tuning, and with a 10-hour fine-tuning which is long enough. As is shown in Fig. 7, the results without fine-tuning have blurred in the details. After 10 minutes of fine-tuning, such artifacts were removed. We also found that further fine-tuning only turns out slight im-

Method	best	second-best	FPS↑
	Time↓		
Neural Body	9.6 hours	0.34	
NeuralVolumes	6 hours	2.3	
ST-NeRF	12 hours	0.04	
iButter	20 mins	3.54	
Ours	2 hours	100	

Table 2. **Quantitative comparison against several methods in terms of training and rendering speed.** Compared with NeuralBody, Neural Volumes, ST-NeRF and iButter, our approach achieves the best performance in FPS metrics and the second best in training or fine-tuning time. For the methods like iButter that use additional fine-tuning do not contain scene-specific training.

provement. Quantitative results in Tab. 3 shows the same phenomenon.

4D Representation in Frequency Domain. We also evaluate the efficiencies of our 4D Representation variants from the perspectives of rendering quality, storage, and consumed time. We set a limit of memory usage as 24GB in our experiments.

As is shown in Fig. 8 and Tab. 4, when the memory is limited, the use of DFT can significantly improve the quality of the results to the case where there is no memory limitation. Also, our model uses the least storage and training time to enable dynamic real-time rendering compared to other methods.




Figure 7. **Qualitative evaluation on time of fine-tuning.** 10-minute fine-tuning achieves the Satisfactory rendering quality while longer fine-tuning does not result in a significant improvement.

Method	best	second-best	MAE↓	LPIPS↓
	PSNR↑	SSIM↑		
w/o fine-tuning	26.02	0.9671	0.0126	0.0678
10-min fine-tuning (ours)	32.93	0.9766	0.0050	0.0340
10-hour fine-tuning	33.39	0.9763	0.0046	0.0336

Table 3. **Qualitative evaluation on time of fine-tuning.** 10-minute fine-tuning achieves the Satisfactory rendering quality while 10-hour fine-tuning does not result in a significant improvement.

6. Discussion

Limitations and Future Work. As the first trial to enable fast generation of octree-based representations and real-time rendering for both static and dynamic scenes, our approach has some limitations.

First, we need dense inputs for static or dynamic scenes. The capturing settings are still expensive and hard to construct. Second, compared with implicit representations such as MLP-based representation for static or dynamic scenes, Fourier PlenOctree still needs larger storage and GPU memory requirement. Also, when the motion of the dynamic scene is faster or the frame number of the multi-view videos is more extended, a higher dimension of Fourier coefficients is needed to keep a high quality of rendered results, which requires higher storage and GPU memory. Third, generalization ability of our generalized nerf may decrease when encountering some remotely exotic objects. Fourth, under the extreme situation of large movements like walking performers on the street, FPO rendering interaction maybe not fluent. Our approach is inefficient as we use the union of visual hulls to initialize Fourier PlenOctrees. But it could be redeemed by using a 3D optical flow method to warp the large movement back to canonical space and then construct FPO. Finally, we are interested in applying our Fourier representation to other neural radiance field representations like [33, 68] for dynamic scenes.

Conclusion. We have presented a novel Fourier PlenOctree (FPO) technique for efficient neural modeling and real-time rendering of dynamic scenes captured under the free-

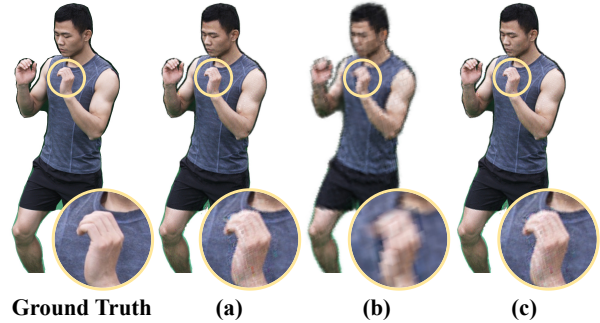


Figure 8. **Qualitative evaluation on DFT.** (a) Our model, using Fourier representation, limited memory w/ DFT, limited memory (ours), (b) w/o DFT, limited memory, (c) w/o DFT, unlimited memory. (b) and (c) are the collection of 60 standard PlenOctrees (different tree structures) without DFT coefficients. Our model is able to use minimal storage with the least training time to achieve high fidelity results. The use of DFT can significantly improve the quality of the results.

Method	best	second-best	Realtime Rendering	Fine-tuning time↓
	PSNR↑	GB↓		
(a)	32.15	7.033	✓	2 hours
(b)	25.97	8.669	✓	10 hours
(c)	32.39	74.32	✗	19 hours

Table 4. **Quantitative evaluation on our Fourier PlenOctree.** (a) Our model, using Fourier representation, limited memory w/ DFT (ours), (b) w/o DFT, limited memory, (c) w/o DFT, unlimited memory. Our model is able to use minimal storage with the least fine-tuning time to achieve high fidelity results.

view video setting. Our coarse-to-fine fusion scheme combines generalizable NeRF with PlenOctree for efficient neural scene construction in minutes. We construct the FPO by tailoring the PlenOctree structure and implicit network to model Fourier coefficients, achieving high-quality rendering for dynamic objects in real-time with compact memory overload. Our experimental results demonstrate the effectiveness and efficiency of FPO for high-quality dynamic scene modeling. With the unique fast generation and real-time rendering ability for dynamic scenes, we believe that our approach serve as a critical step for neural scene modeling, with various potential applications in VR/AR and immersive telepresence.

7. Acknowledgements

This work was supported by Shanghai YangFan Program (21YF1429500), Shanghai Local college capacity building program (22010502800), NSFC programs (61976138, 61977047), the National Key Research and Development Program (2018YFB2100500), STCSM (2015F0203-000-06), and SHMEC (2019-01-07-00-01-E00003). We thank Haizhao Dai for his assistance in scientific visualization.

References

- [1] Naveed Ahmed, Christian Theobalt, Petar Dobrev, Hans-Peter Seidel, and Sebastian Thrun. Robust fusion of dynamic shape and normal capture for high-quality reconstruction of time-varying geometry. In *CVPR*. IEEE Computer Society, 2008. 2
- [2] Benjamin Attal, Eliot Laidlaw, Aaron Gokaslan, Changil Kim, Christian Richardt, James Tompkin, and Matthew O’Toole. Törf: Time-of-flight radiance fields for dynamic scene view synthesis. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 3
- [3] Bruce Guenther Baumgart. *Geometric modeling for computer vision*. Stanford University, 1974. 4
- [4] Mojtaba Bermana, Karol Myszkowski, Hans-Peter Seidel, and Tobias Ritschel. X-fields: Implicit neural view-, light- and time-image interpolation. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2020)*, 39(6), 2020. 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/CVF International Conference on Computer Vision (ICCV)*, pages 14124–14133, October 2021. 3
- [6] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (TOG)*, 34(4):69, 2015. 2
- [7] Mingsong Dou, Philip Davidson, Sean Ryan Fanello, Sameh Khamis, Adarsh Kowdle, Christoph Rhemann, Vladimir Tankovich, and Shahram Izadi. Motion2fusion: Real-time volumetric performance capture. *ACM Trans. Graph.*, 36(6):246:1–246:16, Nov. 2017. 2
- [8] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, Pushmeet Kohli, Vladimir Tankovich, and Shahram Izadi. Fusion4d: Real-time performance capture of challenging scenes. *ACM Trans. Graph.*, 35(4), July 2016. 2
- [9] Robert A Drebin, Loren Carpenter, and Pat Hanrahan. Volume rendering. *ACM Siggraph Computer Graphics*, 22(4):65–74, 1988. 2
- [10] Yilun Du, Yanan Zhang, Hong-Xing Yu, Joshua B. Tenenbaum, and Jiajun Wu. Neural radiance flow for 4d view synthesis and video processing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [11] 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*, 2021. 2
- [12] Peng Gao, Zhengkai Jiang, Haoxuan You, Pan Lu, Steven CH Hoi, Xiaogang Wang, and Hongsheng Li. Dynamic fusion with intra-and inter-modality attention flow for visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6639–6648, 2019. 4
- [13] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14346–14355, October 2021. 1, 3
- [14] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH ’96*, page 43–54, New York, NY, USA, 1996. Association for Computing Machinery. 2
- [15] Yannan He, Anqi Pang, Xin Chen, Han Liang, Minye Wu, Yuexin Ma, and Lan Xu. Challengcap: Monocular 3d capture of challenging human performances using multi-modal references. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11400–11411, June 2021. 2
- [16] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5875–5884, October 2021. 1, 3
- [17] Yuheng Jiang, Suyi Jiang, Guoxing Sun, Zhuo Su, Kaiwen Guo, Minye Wu, Jingyi Yu, and Lan Xu. Neuralfusion: Neural volumetric rendering under human-object interactions. *arXiv preprint arXiv:2202.12825*, 2022. 2
- [18] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8320–8329, 2018. 2
- [19] Hansung Kim, Jean-Yves Guillemaut, Takeshi Takai, Muhammad Sarim, and Adrian Hilton. Outdoor dynamic 3d scene reconstruction. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(11):1611–1622, 2012. 2
- [20] Naruya Kondo, Yuya Ikeda, Andrea Tagliasacchi, Yutaka Matsuo, Yoichi Ochiai, and Shixiang Shane Gu. Vaxnerf: Revisiting the classic for voxel-accelerated neural radiance field, 2021. 3
- [21] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *International journal of computer vision*, 38(3):199–218, 2000. 2
- [22] Marc Levoy and Pat Hanrahan. Light field rendering. In *SIGGRAPH*, pages 31–42, 1996. 2
- [23] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T. Freeman. Learning the depths of moving people by watching frozen people. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4516–4525, 2019. 2
- [24] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6498–6508, June 2021. 2
- [25] David B Lindell, Julien NP Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural volume rendering. In *CVPR*, 2021. 3
- [26] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 2020. 3

- [27] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Trans. Graph.(ACM SIGGRAPH Asia)*, 2021. [2](#)
- [28] 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, 2019. [1](#), [6](#)
- [29] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM Trans. Graph.*, 40(4), jul 2021. [1](#)
- [30] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (TOG)*, 39(4):71–1, 2020. [2](#)
- [31] Zhaoyang Lv, Kihwan Kim, Alejandro Troccoli, Deqing Sun, James M. Rehg, and Jan Kautz. Learning rigidity in dynamic scenes with a moving camera for 3d motion field estimation. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 484–501, Cham, 2018. Springer International Publishing. [2](#)
- [32] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020. [1](#), [2](#), [6](#)
- [33] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022. [3](#), [8](#)
- [34] Armin Mustafa, Hansung Kim, Jean-Yves Guillemaut, and Adrian Hilton. Temporally coherent 4d reconstruction of complex dynamic scenes. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4660–4669, 2016. [2](#)
- [35] Richard A. Newcombe, Dieter Fox, and Steven M. Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 343–352, 2015. [2](#)
- [36] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE International Symposium on Mixed and Augmented Reality*, pages 127–136, 2011. [2](#), [4](#)
- [37] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. 2020. [2](#), [3](#)
- [38] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5865–5874, October 2021. [2](#)
- [39] 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 Trans. Graph.*, 40(6), dec 2021. [2](#), [3](#)
- [40] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, 2021. [2](#)
- [41] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. [2](#), [6](#)
- [42] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)*, 36(6):1–11, 2017. [2](#)
- [43] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. [2](#)
- [44] René Ranftl, Vibhav Vineet, Qifeng Chen, and Vladlen Koltun. Dense monocular depth estimation in complex dynamic scenes. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4058–4066, 2016. [2](#)
- [45] 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/CVF International Conference on Computer Vision (ICCV)*, pages 14335–14345, October 2021. [1](#), [3](#)
- [46] Chris Russell, Rui Yu, and Lourdes Agapito. Video pop-up: Monocular 3d reconstruction of dynamic scenes. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 583–598, Cham, 2014. Springer International Publishing. [2](#)
- [47] Steven M Seitz and Charles R Dyer. Photorealistic scene reconstruction by voxel coloring. *International Journal of Computer Vision*, 35(2):151–173, 1999. [2](#)
- [48] Vincent Sitzmann, Semon Rezchikov, William T. Freeman, Joshua B. Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. In *Proc. NeurIPS*, 2021. [2](#)
- [49] Zhuo Su, Lan Xu, Zerong Zheng, Tao Yu, Yebin Liu, and Lu Fang. Robustfusion: Human volumetric capture with data-driven visual cues using a rgbd camera. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 246–264, Cham, 2020. Springer International Publishing. [2](#)
- [50] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. *arXiv preprint arXiv:2111.11215*, 2021. [3](#)
- [51] Richard Szeliski and Polina Golland. Stereo matching with transparency and matting. In *Sixth International Conference on Computer Vision (IEEE Cat. No. 98CH36271)*, pages 517–524. IEEE, 1998. [2](#)
- [52] Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P. Srinivasan, Jonathan T. Barron, and Ren

- Ng. Learned initializations for optimizing coordinate-based neural representations. In *CVPR*, 2021. 3
- [53] Aparna Taneja, Luca Ballan, and Marc Pollefeys. Modeling dynamic scenes recorded with freely moving cameras. In Ron Kimmel, Reinhard Klette, and Akihiro Sugimoto, editors, *Computer Vision – ACCV 2010*, pages 613–626, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. 2
- [54] Edgar Treitsch, 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/CVF International Conference on Computer Vision (ICCV)*, pages 12959–12970, October 2021. 2
- [55] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d representation and rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15182–15192, 2021. 3
- [56] Daniel Vlasic, Pieter Peers, Ilya Baran, Paul Debevec, Jovan Popović, Szymon Rusinkiewicz, and Wojciech Matusik. Dynamic shape capture using multi-view photometric stereo. In *ACM SIGGRAPH Asia 2009 Papers*, SIGGRAPH Asia '09, New York, NY, USA, 2009. Association for Computing Machinery. 2
- [57] Michael Waechter, Nils Moehrl, and Michael Goesele. Let there be color! large-scale texturing of 3d reconstructions. In *European conference on computer vision*, pages 836–850. Springer, 2014. 2
- [58] Liao Wang, Ziyu Wang, Pei Lin, Yuheng Jiang, Xin Suo, Minye Wu, Lan Xu, and Jingyi Yu. ibutter: Neural interactive bullet time generator for human free-viewpoint rendering. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4641–4650, 2021. 2, 3, 6
- [59] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. IBRNet: Learning multi-view image-based rendering. In *CVPR*, 2021. 2, 3
- [60] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. NeX: Real-time view synthesis with neural basis expansion. In 2021. 3
- [61] Daniel N Wood, Daniel I Azuma, Ken Aldinger, Brian Curless, Tom Duchamp, David H Salesin, and Werner Stuetzle. Surface light fields for 3d photography. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 287–296, 2000. 2
- [62] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9421–9431, June 2021. 2
- [63] Lan Xu, Wei Cheng, Kaiwen Guo, Lei Han, Yebin Liu, and Lu Fang. Flyfusion: Realtime dynamic scene reconstruction using a flying depth camera. *IEEE Transactions on Visualization and Computer Graphics*, 27(1):68–82, 2021. 2
- [64] Lan Xu, Yebin Liu, Wei Cheng, Kaiwen Guo, Guyue Zhou, Qionghai Dai, and Lu Fang. Flycap: Markerless motion capture using multiple autonomous flying cameras. *IEEE Transactions on Visualization and Computer Graphics*, 24(8):2284–2297, Aug 2018. 2
- [65] Lan Xu, Zhuo Su, Lei Han, Tao Yu, Yebin Liu, and Lu Fang. Unstructuredfusion: Realtime 4d geometry and texture reconstruction using commercial rgbd cameras. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(10):2508–2522, Oct. 2020. 2
- [66] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields, 2022. 3
- [67] 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. June 2020. 2
- [68] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. *arXiv preprint arXiv:2112.05131*, 2021. 3, 8
- [69] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021. 1, 3, 5
- [70] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021. 3
- [71] Jiakai Zhang, Xinhang Liu, Xinyi Ye, Fuqiang Zhao, Yanshun Zhang, Minye Wu, Yingliang Zhang, Lan Xu, and Jingyi Yu. Editable free-viewpoint video using a layered neural representation. *ACM Transactions on Graphics (TOG)*, 40(4):1–18, 2021. 2, 6
- [72] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 6
- [73] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. Humannerf: Efficiently generated human radiance field from sparse inputs. *arXiv preprint arXiv:2112.02789*, 2021. 2