

# Frequency-Modulated Point Cloud Rendering with Easy Editing

Yi Zhang<sup>1\*</sup> Xiaoyang Huang<sup>1\*</sup> Bingbing Ni<sup>1†</sup> Teng Li<sup>2</sup> Wenjun Zhang<sup>1</sup>  
<sup>1</sup>Shanghai Jiao Tong University, Shanghai 200240, China <sup>2</sup>Anhui University  
 {yizhangphd, huangxiaoyang, nibingbing}@sjtu.edu.cn




Figure 1. We propose a novel frequency-modulated point cloud rendering pipeline, which enables high fidelity local detail reconstruction, real-time rendering and user-friendly editing. Specifically, we can render sharper textures with  $1700\times$  faster than NeRF [21] and support both object-level editing and scene composition.

## Abstract

We develop an effective point cloud rendering pipeline for novel view synthesis, which enables high fidelity local detail reconstruction, real-time rendering and user-friendly editing. In the heart of our pipeline is an adaptive frequency modulation module called Adaptive Frequency Net (AFNet), which utilizes a hypernetwork to learn the local texture frequency encoding that is consecutively injected into adaptive frequency activation layers to modulate the implicit radiance signal. This mechanism improves the frequency expressive ability of the network with richer frequency basis support, only at a small computational budget. To further boost performance, a preprocessing module is also proposed for point cloud geometry optimization via point opacity estimation. In contrast to implicit rendering, our pipeline supports high-fidelity interactive editing based on point cloud manipulation. Extensive experimental results on NeRF-Synthetic, ScanNet, DTU and Tanks and Temples datasets demonstrate the superior performances achieved by our method in terms of PSNR, SSIM and LPIPS, in comparison to the state-of-the-art. Code is released at <https://github.com/yizhangphd/FreqPCR>.

## 1. Introduction

Photo-realistic rendering and editing of 3D representations is a key problem in 3D computer vision and graphics with numerous applications, such as computer games,

VR/AR, and video creation. In particular, recently introduced neural radiance field (NeRF) [21] has inspired some follow-up works aiming to editable rendering [15, 17, 20, 47, 49, 55]. Due to the deeply coupled black-box network, NeRF-based object-level editing usually requires a pre-trained segmentation model to separate the objects to be edited [17, 55]. Although some recent voxel-based variants of NeRF [47, 58] achieve multi-scene composition, they still lack the ability to extract target objects from voxels.

In contrast to implicit rendering, point cloud rendering [1, 6, 13, 18, 33, 36, 57] is a promising editable rendering model. On the one hand, explicit 3D representations are better for interactive editing. On the other hand, the geometric priors provided by point clouds can help us avoid massive sampling in volume rendering methods, which can meet the requirements of some real-time applications. As a class of representative point cloud rendering methods, NPBG and NPBG++ [1, 33] achieve real-time rendering by using point-wise features for encoding appearance information and an U-Net [35] for decoding, respectively. However, the parameter quantity increases with the size of point cloud, which may limit their application due to the excessive computational and memory complexity. Huang *et al.* [13] combine point clouds with implicit rendering, where explicit point clouds are used to estimate the geometry, and implicit radiance mapping is used to predict view-dependent appearance of surfaces. However, quantitative evaluation of its rendering results is significantly lower than that of implicit rendering methods such as NeRF [21], mainly due to the following reasons: 1) the color of each viewing ray only depends on

\*Equal contribution.

†Corresponding author: Bingbing Ni.

a single surface point, thus without multiple sample color aggregation for error attenuation, surface based rendering techniques require radiance mapping to have a more precise and expressive frequency encoding ability; and 2) defects of the point cloud geometry reconstructed by MVSNet [56] cause wrong surface estimation. To this end, we introduce Adaptive Frequency Net (AFNet) to improve frequency expression ability of the radiance mapping and a preprocessing module for point cloud geometry optimization.

Radiance mapping, also known as radiance field, is a type of Implicit Neural Representation (INR). There have been some studies [2, 7, 32, 41, 45, 46, 60] on the expressive power and inductive bias of INRs. The standard Multi-layer Perceptrons (MLPs) with ReLU activation function are well known for the strong spectral bias towards reconstructing low frequency signals [32]. Some recent works [7, 41, 46] introduce strategies to enhance the high-frequency representation of MLPs from a global perspective. However, from a local perspective, the frequencies of a 3D scene are region-dependent and most real objects are composed by both weak textured regions and strong textured ones. Motivated by this, we design a novel adaptive frequency modulation mechanism based on HyperNetwork architecture [11], which learns the local texture frequency and injects it into adaptive frequency activation layers to modulate the implicit radiance signal. The proposed mechanism can predict suitable frequency without frequency supervision and modulate the radiance signal with adaptive frequency basis support to express more complex textures at negligible computational overhead.

Previous surface point-based works [1, 13, 33] could not optimize the point cloud geometry because they keep only the closest point as a surface estimation for each ray. But if we sample multiple points per ray during rendering, it will greatly reduce the rendering speed. Therefore, we use the volume rendering method as a preprocessing module to optimize the point cloud geometry. Specifically, we keep more points in the pixel buffer and learn the opacity of each point based on volume rendering. We find in our experiments that for some poorly reconstructed scenes, point cloud geometry optimization can improve the rendering PSNR by 2-4dB and avoid rendering artifacts.

For rigid object editing, we follow the deformation field construction [28–31, 48] to render the edited point cloud. Point cloud can be seen as a bridge between user editing and deformation field to achieve interactive editing and rendering. Users only need to edit the point cloud, and the deformation field between the original 3D space and the deformed space is easy to obtain by querying the corresponding transformations performed on point cloud. Moreover, to avoid cross-scene training in multi-scene composition application, we develop a masking strategy based on depth buffer to combine multiple scenes in pixel level.

We evaluate our method on NeRF-Synthetic [21], ScanNet [5], DTU [14] and Tanks and Temples [16] datasets and comprehensively compare the proposed method with other works in terms of performance (including PSNR, SSIM and LPIPS), model complexity, rendering speed, and editing ability. Our performance outperforms NeRF [21] and all surface point-based rendering methods [1, 13, 33], and is comparable to Compressible-composable NeRF (CC-NeRF), *i.e.*, the latest NeRF-based editable variant. We achieve a real-time rendering speed of 39.27 FPS on NeRF-Synthetic, which is 1700× faster than NeRF and 37× faster than CCNeRF. We also reproduce the scene editing of Object-NeRF [55] and CCNeRF [47] on ToyDesk [55] and NeRF-Synthetic [21] dataset, respectively. As shown in Fig. 1, we achieve real-time rendering with sharper details and user-friendly editing. The above results demonstrate that our method is comprehensive in terms of both rendering and editing and has great application potentials.

## 2. Related Work

**Implicit Neural Representations.** The goal of INRs is to encode a continuous target signal using a neural network, by representing the mapping between input coordinates and signal values. This technique has been used to represent various objects such as images [23, 43], shapes [4, 9, 27], scenes [21, 42], and textures [12, 25]. Tancik *et al.* [46] propose Fourier Feature Network (FFN), showing that feeding input coordinates through a Fourier feature mapping enables MLP to learn high-frequency functions in low dimensional domains. Sitzmann *et al.* [41] propose to use MLP with sinusoidal activations. Fathony *et al.* propose MFN [7], which employs a Hadamard product between linear layers and nonlinear activation functions. In [60], Yüce *et al.* prove that FFN [46], SIREN [41] and MFN [7] have the same expressive power as a structured signal dictionary.

**3D Representations for View Synthesis.** Implicit representations often use neural networks to encode 3D scenes [21, 24, 26, 27, 39, 50, 51, 61]. NeRF [21] uses a 5D function to represent the scene and applies volumetric rendering for view synthesis. Many recent works combine NeRF with some explicit representation to improve rendering quality or speed. NSVF [19] uses a sparse voxel octree to achieve fast rendering. Yu *et al.* propose PlenOc-trees [58], a data structure derived from NeRF [21] which enables highly efficient rendering. Plenoxels [8] represents a scene as a sparse 3D grid with spherical harmonics. Xu *et al.* [52] propose Point-NeRF, which is a localized neural representation, combining volumetric radiance fields with point clouds. Instant-NGP [22] uses multi-resolution hashing for efficient encoding and also leads to high compactness. TensorRF [3] factorizes the scene tensor into lower-rank components. Huang *et al.* [13] com-

bine point cloud with implicit radiance mapping to achieve the state-of-the-art performance of point cloud rendering. However, their performance is still worse than NeRF [21]. In this work, we greatly improve the performance of point cloud rendering by introducing AFNet and a preprocessing module for geometry optimization.

**Editable Rendering.** Editing ability is important for 3D representations. Some recent works [10, 44, 53–55, 59] improve NeRF to editable version. Yang *et al.* [55] propose Object-NeRF for editable rendering, while it requires a 2D segmentation model to extract the target object and renders at a slow speed. Yuan *et al.* propose NeRF-Editing [59], which allows users to perform controllable shape deformation on the implicit representation. Yang *et al.* [54] present a mesh-based representation by encoding the implicit field with disentangled geometry and texture codes on mesh vertices, which facilitates a set of editing functionalities. Tang *et al.* propose CCNeRF [47], which could transform and compose several scenes into one scene with a small model.

### 3. Method

In this section, we first revisit the state-of-the-art point cloud rendering technique and discuss its dilemma of insufficient frequency representation, in Sec. 3.1. In Sec. 3.2, we present our proposed frequency-modulated radiance mapping technique to address above limitation, followed by the introduction to our end-to-end point cloud rendering pipeline in Sec. 3.3. We also introduce a novel point cloud geometry optimization pipeline as an optional preprocessing module in Sec. 3.4. Applications of the above pipeline on interactive object level editing and multi-scene composition are presented in Sec. 3.5.

#### 3.1. Revisiting Point Cloud Rendering

Some traditional point cloud rendering methods focus on fast rendering of large-scale point clouds [37, 38], while they are not as photorealistic as neural rendering (*e.g.*, not smooth and lack of view-dependent appearance). Different from traditional works, our approach builds on a state-of-the-art point-based neural rendering method [13], which utilizes point cloud to determine the position of the surface by the following coordinate rectification

$$\mathbf{x} = \mathbf{o} + \frac{z}{\cos \theta} \times \mathbf{d}, \quad (1)$$

where  $\mathbf{o}$ ,  $\mathbf{d}$  and  $\theta$  denote camera position, normalized ray direction and angle between the ray and the  $z$ -axis pointing to the center of the scene, respectively. In practice,  $z$  value is obtained by querying the depth buffer during rasterization. As point clouds are discrete, each point is expanded to a disk of radius  $r$  during rasterization [63]. The estimated surface coordinates  $\mathbf{x}$  together with ray direction  $\mathbf{d}$  are fed

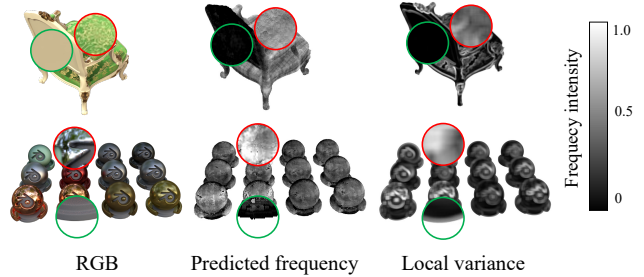


Figure 2. Original RGB images, frequency maps predicted by our hypernetwork, and local variance of original images for reference. We find that the frequencies are large for strongly textured regions (shown as gray and white color), and the frequencies are close to zeros for weakly textured regions (shown as black color).

to an MLP  $\mathbf{F}_\Theta$  ( $\Theta$  is the network parameter set) to obtain the features  $\mathbf{L}$  which encode the color of the surface point

$$\mathbf{L} = \mathbf{F}_\Theta(\mathbf{x}, \mathbf{d}). \quad (2)$$

Compared to volume rendering-based models [21, 52], this method avoids heavy sampling by estimating surface coordinates from point clouds when calculating ray color, and thus achieving real-time rendering. As point clouds are usually noisy and the initial depth obtained by rasterization might also be inaccurate, previous methods [1, 13, 33] often use U-Net structure [35] for further refinement.

**Limitations.** The limitations of the above surface point-based rendering scheme are obvious. Namely, the point cloud quality is usually poor due to sparsity and large number of outliers, which leads to lower-fidelity rendering especially in high frequency shapes, *i.e.*, highly-3D-textured regions. Simply increasing the network can not solve the problem. On the contrary, prevailing volumetric rendering techniques such as NeRF [21] do not suffer from this issue (*i.e.*, usually with higher rendering quality), since through each sampled ray there exist multiple sampling points near the surface fused together to contribute the target color, which is more stable and sufficient for dealing with high frequency shapes. Therefore, it is highly demanded to enhance the frequency expression ability of point cloud based implicit surface radiance mapping method at a relatively low computational cost overhead.

#### 3.2. Frequency-Modulated Radiance Mapping

**Motivation.** To achieve radiance mapping, also known as radiance field, previous works input a spatial coordinate and infer the color or other features associated with this spatial position via standard MLP with ReLU activation layer. However, this classical neural network is well-known for its strong spectral bias towards generating lower frequency patterns [32], which has obvious difficulty in handling high frequency shape details, *e.g.*, 3D textures of objects and occupancy of edges. Recent works [7, 41, 46] analyze the expres-

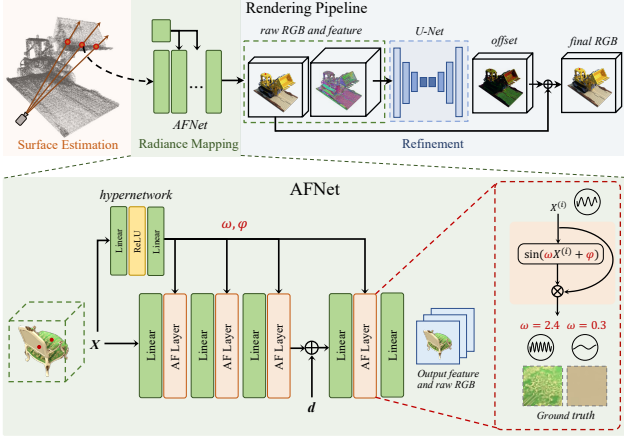


Figure 3. Overview of our rendering pipeline. In the heart of the pipeline is an adaptive frequency modulation module called AFNet, which utilizes a hypernetwork to learn local texture frequency encoding that is consecutively injected to adaptive frequency activation layers to modulate the implicit radiance signal.

side power and inductive bias of the network and enhance the fitting ability of the network from a global perspective, while we emphasize the local perspective of the problem, *i.e.*, the frequencies of a 3D scene are region-dependent. Namely, most real objects are composed by both weak textured regions and strong textured ones. For example, in the Chair scene of NeRF-Synthetic dataset, the back of the chair belongs to weak texture areas, and the cushion with patterns belongs to strong texture areas, as shown in Fig. 2. If the signal frequency does not match the target area, the implicit radiance mapping would be poor due to misaligned frequency functional support. Thus, the above observations motivate us to propose an Adaptive Frequency Net (AFNet) for frequency-modulated radiance mapping, as shown in Fig. 3, which is capable of selecting high frequency primitives for modeling rich-3D-textured regions while uses low frequency basis for presenting non-textured regions, in order to yield overall high-fidelity rendering quality.

**Adaptive Frequency Estimation and Modulation.** The major working mechanism of the proposed AFNet is to first estimate the local frequency information for each point, and then this frequency encoding is injected to the proposed Adaptive Frequency Activation Layers (AF Layers) which adaptively select a proper set of frequency component, providing a more precise support for generating high-fidelity reconstruction, as shown in Fig. 3. Therefore, the core building block of the proposed AFNet includes 1) AF Layers and 2) a hypernetwork for local frequency estimation, which are introduced in detail as follows.

**Adaptive Frequency Activation Layer.** Assuming local frequency information is successfully estimated and encoded as scalar frequencies and phases, our AF layers first encode the input radiance signal with a nonlinear filter (we

use a sinusoidal filter) parameterized by the obtained point-wise adaptive frequency and phase. Then, a Hadamard product between the filtered feature and the input feature is taken as the activated output, which could be regarded as a frequency modulation on the input signal. Mathematically, suppose that the input feature of  $i$ -th layer is  $\mathbf{X}^{(i)} \in \mathbb{R}^{M \times C}$ , in which  $M$  and  $C$  are point number and feature dimension respectively, and the adaptive frequencies and phases are  $\omega^{(i)}, \phi^{(i)} \in \mathbb{R}^{M \times 1}$ , whose dimensions will be broadcast to  $\mathbb{R}^{M \times C}$  during computation. Then the adaptive frequency activation layers can be expressed as:

$$\mathbf{Y}^{(i)} = \mathbf{X}^{(i)} \circ \sin(\omega^{(i)} \circ \mathbf{X}^{(i)} + \phi^{(i)}). \quad (3)$$

The design of Eqn. 3 has a theoretical justification. We measure the local frequency by the derivative of the color with respect to the spatial coordinate, and adaptively adjust the derivative by multiplying the feature  $\mathbf{X}^{(i)}$  by the predicted frequency  $\omega^{(i)}$ . We then use the sine function to constrain the numerical range. The final form,  $\mathbf{X}^{(i)} \circ \sin(\cdot)$ , is borrowed from MFN [7].

**Hypernetwork for Frequency Estimation.** The adaptive frequencies and phases are expected to contain important information of the 3D texture variation on the surface position; however, to obtain the accurate frequency representation of each point is non-trivial. Inspired by HyperNetwork [11], an approach of using a small network to generate the weights for a larger network, we use a two-layer MLP as hypernetwork to infer the frequency and phase parameters from each spatial point. In Fig. 2, we show the predicted frequency of the 4-th adaptive frequency activation layer. As seen in the RGB images, the back of the chair and the base of the ball have weak texture, while the frequency predicted by hypernetwork is closed to 0 (shown as black color). For strongly textured regions, such as patterned cushions in chair and the metal material, the hypernetwork predicts large frequency (shown as grey and white). These observations empirically indicates that the proposed hypernetwork is capable of learning correct frequency information even there is **NO** supervision on frequency label.

### 3.3. End-to-end Rendering Pipeline

Our pipeline consists of three stages: surface estimation, radiance mapping and refinement, as shown in Fig. 3. In the surface estimation stage, we first rasterize the point cloud  $\mathbf{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\} \in \mathbb{R}^{N \times 3}$  to obtain the depth buffer  $\mathbf{z} \in \mathbb{R}^{H \times W}$  at a given view. Specifically, we transform the point cloud from the world coordinate space to the normalized device coordinate space, in which  $\mathbf{p}_k \in [-1, 1] \times [-1, 1] \times [z_{near}, z_{far}]$ ,  $k = 1, 2, \dots, N$ . Similarly, we transform the pixel coordinates  $\mathbf{I} \in \mathbb{R}^{H \times W \times 2}$  to normalized space such that  $\mathbf{I}_{i,j} \in [-1, 1]^2$ . Then, each point  $\mathbf{p}_k$  is expanded to a disk of radius  $r_k$  in  $xOy$  plane.



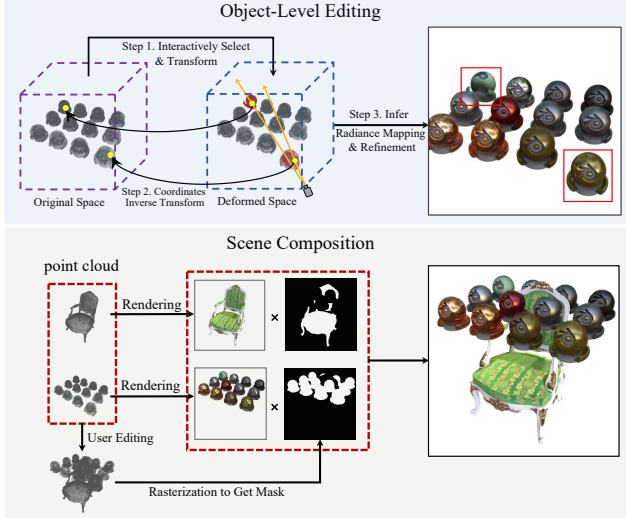


Figure 4. Object-level editing and scene composition pipelines.

The depth buffer  $z_{i,j}$  is the depth of the nearest disk covering pixel  $\mathbf{I}_{i,j}$ , that is

$$z_{i,j} = \min_k \mathbf{p}_k^z, \quad s.t. (\mathbf{p}_k^x - \mathbf{I}_{i,j}^x)^2 + (\mathbf{p}_k^y - \mathbf{I}_{i,j}^y)^2 < r_k^2, \quad (4)$$

which can be seen as a special case of surface splatting [63]. Then, we use Eqn. 1 to estimate the coordinate of the surface where each ray intersects.

In radiance mapping stage, the estimated surface coordinates are fed into the proposed AFNet to predict the raw RGB values and feature vectors. In the refinement stage, the feature vectors will be fed into an U-Net to predict the offsets. The final RGBs are obtained by adding the raw RGB values and offsets. AFNet and U-Net are trained with ground truth images as supervision by using L2 loss and perceptual loss [62].

### 3.4. Preprocessing: Geometry Optimization

In the meantime, point clouds obtained by MVS method [56] or 3D scans inevitably contain noise or missing parts of the surface, which harms the rendering quality or even produces artifacts. In view of this, we design a point cloud geometry optimization strategy consisting of a point cloud denoising step and a completion step, considered as an optional preprocessing module in our framework, to further enhance rendering quality.

**Denoising.** We design a volume rendering-based pipeline specifically to optimize the point cloud geometry. In rasterization, we maintain the depth buffer of  $k$  points instead of the nearest point. The depth values are used to compute the rectified coordinates of the points using Eqn. 1. The rectified coordinates, combined with ray directions, are fed into a MLP to estimate colors. Besides, an optimizable opacity parameter  $\alpha_i$  is assigned to each point  $\mathbf{p}_i$ . Pixel RGB values are obtained using the volume rendering equation. Inspired

by Point-NeRF [52], we train the MLP and the opacity parameters to fit the ground-truth images with the following sparse regularization:

$$\mathcal{L}_{\text{sparse}} = \frac{1}{N} \sum_{i=1}^N [\log(\alpha_i) + \log(1 - \alpha_i)]. \quad (5)$$

When the training converges, we remove those low-opacity points with a threshold  $\mu$ .

**Completion.** For each view in training set, if a pixel in the ground truth image is occupied, but the buffer obtained by rasterization is empty, we add a set of points with opacity parameters along the pixel ray in the point clouds. To determine the depth range of the added points, we refer to the depth range of adjacent pixel buffer. Although the completion step introduces additional noise, it would be removed in the denoising stage of the next iteration.

Above denoising and completion steps run alternatively until convergence. Please see supplementary material for more details.

### 3.5. Scene Editing

In this section, we introduce an editing method based on the proposed rendering pipeline. At object level, we leverage point clouds to bridge the gap between user editing and deformation field. At scene level, the proposed method can combine different scenes without cross-scene training.

**Object-level Editing.** As an explicit 3D representation, point clouds are ideal for interactive editing. After editing target points, we can estimate the new surface coordinates of the edited point cloud by rasterization. Since radiance mapping encodes appearance of the surface in original 3D space, we need to transform the estimated surface points to original encoding space, which can be regarded as a deformation of the original space, *i.e.*, deformation field [28, 31]. Specifically, we keep the indices of the points in pixel buffer during rasterization and query the transformations performed on those points. Then we perform the corresponding inverse transformation on the estimated surface coordinates to obtain the locations in original 3D space, which replaces the calculation of deformation field, as shown in Fig. 4. Finally, we feed the transformed surface points to radiance mapping module followed by refinement module to obtain the rendering result.

**Scene Composition.** A direct way to render a combined point cloud is to use respective AFNets for radiance mapping, and a shared U-Net for refinement. However, due to the inductive bias of U-Net, the texture of different scenes will affect each other, which would reduce the rendering quality. Moreover, cross-scene training of U-Net impairs the flexibility of editing, *i.e.*, adding a new scene to an existing composition requires retraining together. To solve the above problems, we compose the scenes in pixel level. As

Method	NeRF-Synthetic			Tanks&Temples			ScanNet			DTU		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
NeRF [21]	31.01	0.947	0.081	25.78	0.864	0.198	25.74	0.780	0.537	26.92	0.909	0.198
CCNeRF-CP [47]	30.55	0.935	0.076	27.01	0.879	0.180	24.65	0.774	0.542	26.79	0.907	0.178
CCNeRF-HY-S [47]	31.22	0.947	0.074	27.53	0.901	0.177	25.17	0.781	0.539	27.23	0.910	0.171
NPBG [1]	28.10	0.923	0.077	25.97	0.889	0.137	25.09	0.737	0.459	26.00	0.895	0.130
NPBG++ [33]	28.12	0.928	0.076	26.04	0.892	0.130	25.27	0.772	0.448	26.08	0.895	0.131
Huang <i>et al.</i> [13]	28.96	0.932	0.061	26.35	0.893	0.130	25.88	0.794	0.415	26.22	0.900	0.132
Ours	<b>31.24</b>	<b>0.950</b>	<b>0.049</b>	<b>27.79</b>	<b>0.902</b>	<b>0.125</b>	<b>26.66</b>	<b>0.803</b>	<b>0.400</b>	<b>27.27</b>	<b>0.908</b>	<b>0.129</b>

Table 1. Quantitative evaluation on NeRF-Synthetic [21], Tanks and Temples [16], ScanNet [5] and DTU [14] datasets. Since our work focuses on surface point-based method, which supports real-time and editable rendering, we mainly compare with NPBG [1], NPBG++ [33] and Huang *et al.* [13]. We also show the performance of NeRF [21] and CCNeRF [47] for reference.

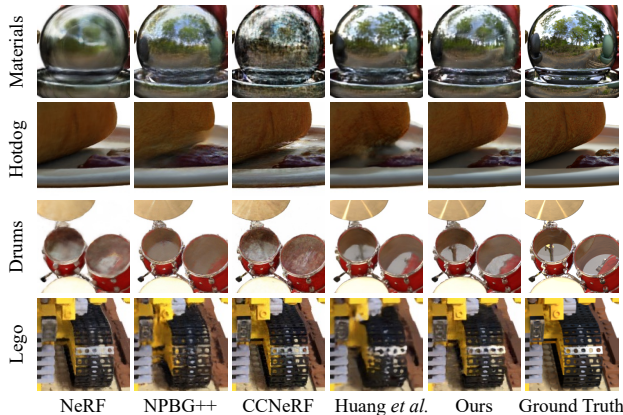


Figure 5. Qualitative evaluations on NeRF-Synthetic [21] dataset. Our method renders sharper detailed textures without artifacts.

shown in Fig. 4, we first rasterize the combined point clouds to determine which scene each pixel belongs to and generate a mask for each scene. Then we render each scene separately and multiply the masks with the rendered images. We combine the masked images to obtain the final result.

## 4. Experiments

### 4.1. Experimental Settings

**Datasets and Compared Methods.** Since our work focuses on surface point-based rendering methods, which enable real-time and editable rendering, we mainly compare with the following methods: 1) NPBG [1]: A famous surface point-based rendering method which enables real-time rendering and editing; 2) NPBG++ [33]: The improved version of NPBG; and 3) Huang *et al.* [13]: The state-of-the-art surface point-based rendering pipeline. We first compare our performance with the above methods on NeRF-Synthetic [21], Tanks and Temples [16], ScanNet [5] and DTU [14] datasets. To further demonstrate the superiority of our method, we conduct a comprehensive comparison with volume rendering-based methods, including NeRF [21] and its variants, on NeRF-Synthetic [21] and Tanks and Temples [16] datasets. Please see supplementary material for more details.




Figure 6. Qualitative evaluations on Tanks and Temples [16] dataset. Our method renders sharper detailed textures.

**Implementation Details.** There are two ways of rasterization: PyTorch3D [34] for headless server and OpenGL [40] for real-time rendering. Since our training is performed on a headless server, our main results are based on PyTorch3D rasterization. The widths of AFNet and hypernetwork are both 256, and camera ray directions are concated to the third layer output feature of AFNet. Moreover, we modify the positional encoding of NeRF [21] to make it more suitable for modulation, which is described in detail in supplementary material. The architecture of our U-Net is consistent with NPBG [1], *i.e.*, gated blocks are used, consisting of gated convolution and instance normalization. The U-Net down-samples the feature map for four times and the dimensions of each layer are 16, 32, 64, 128, 256. We train the AFNet and U-Net with L2 loss and perceptual loss [62] on a single NVIDIA GeForce RTX 3090 GPU. For preprocessing of point cloud optimization, we use only L2 loss and a 5-layers MLP with 256 width, where ray directions are concated to the third layer output feature. We use Adam optimizer for training, with a batch size of 2. The ini-

Method	Size(MB)	FPS	Editing Ability		NeRF-Synthetic			Tanks&Temples		
			Object	Scene	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
NeRF [21]	5.0	0.023	$\times$	$\times$	31.01	0.947	0.081	25.78	0.864	0.198
NSVF [19]	16.0	0.815	$\times$	$\checkmark$	31.75	0.953	0.047	28.40	0.900	0.153
Object-NeRF [55]	121.2	0.1	$\checkmark$	$\times$	31.19	0.949	0.079	25.96	0.866	0.194
PlenOctrees [58]	1976.3	168	$\times$	$\checkmark$	31.71	0.958	0.053	27.99	0.917	0.131
Point-NeRF [52]	20.0	0.125	$\times$	$\times$	33.00	0.978	0.055	29.61	0.954	0.115
Plenoxels [8]	778.1	15	$\times$	$\checkmark$	31.71	0.958	0.049	27.43	0.906	0.142
TensorRF [3]	71.8	1.15	$\times$	$\checkmark$	33.14	0.963	0.056	28.56	0.920	0.118
Instant-NGP [22]	63.3	60	$\times$	$\times$	33.18	0.963	0.050	28.78	0.925	0.113
CCNeRF-CP [47]	4.4	1.05	$\times$	$\checkmark$	30.55	0.935	0.076	27.01	0.879	0.180
CCNeRF-HY-S [47]	68.9	1.05	$\times$	$\checkmark$	31.22	0.947	0.074	27.53	0.901	0.177
NPBG [1]	44.2	33.4	$\checkmark$	$\checkmark$	28.10	0.923	0.077	25.97	0.889	0.137
NPBG++ [33]	28.6	35.4	$\checkmark$	$\checkmark$	28.12	0.928	0.076	26.04	0.892	0.130
Huang <i>et al.</i> [13]	18.5	39.1	$\checkmark$	$\checkmark$	28.96	0.932	0.061	26.35	0.893	0.130
Ours	11.8	39.3	$\checkmark$	$\checkmark$	31.24	0.950	0.049	27.79	0.902	0.125

Table 2. Comprehensive evaluation. We report the model size (including point cloud size for point-based methods), FPS on NeRF-Synthetic [21] dataset with  $800\times 800$  resolution, editing ability of object-level transform and scene composition, and performance on NeRF-Synthetic [21] and Tanks and temples [16] datasets. The results demonstrate the excellent comprehensiveness of our method.

tial learning rates of AFNet and U-Net are  $5e-4$  and  $1.5e-4$  respectively.

## 4.2. Results of View Synthesis

We present quantitative comparison results on three performance metrics (PSNR, SSIM and LPIPS) in Tab. 1. As can be seen, we outperform the previous point-based methods with a significant margin on all datasets. Our method outperforms NeRF [21] and is comparable to the high-profile version of CCNeRF [47]. Figs. 5 and 6 present a qualitative comparison on NeRF-Synthetic [21] and Tanks and Temples [16] dataset. As can be seen, our method produces sharper details without introducing artifacts, especially for highly textured regions.

**Comprehensive Analysis.** To further illustrate the superiority of our method, we expand the quantitative experiments with more comparison works in Tab. 2, and comprehensively compare them from four aspects: model size, rendering speed, editing ability, and performance. It can be seen that almost all the follow-up works of NeRF are difficult to take into account the four aspects at the same time. PlenOctrees [58] achieves the state-of-the-art FPS, but its application scenarios are limited due to its huge model size. Instant-NGP [22] has high performance and real-time rendering speed, but lacks editing abilities. Object-NeRF [55] and CCNeRF [47] can achieve object-level editing and scene composition respectively, but cannot achieve real-time rendering. Our method achieves high performance and flexible editing with a small model size while rendering in real-time.

**Ablation Study.** The frequencies predicted by our hypernetwork are highly correlated with the true frequencies reflected by local variance, as shown in Fig. 7. Note that there

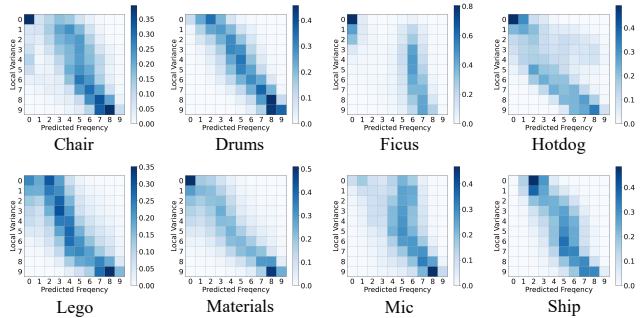


Figure 7. Confusion matrix between predicted frequency map and actual local variance on 200 test views of NeRF-Synthetic dataset. The frequency interval is divided into 10 segments, and from 0 to 9 are low frequency to high frequency. The frequencies predicted by our hypernetwork are highly correlated with the true frequencies reflected by local variance.

is a reasonable gap between our predicted 3D texture frequencies and 2D image frequencies. For example, in Ficus scene, the occlusion between leaves and background produces a high frequency region in 2D image, but the 3D texture of leaves is weak (just green). Moreover, we claim that the frequencies predicted by hypernetwork are adaptive, which can modulate the radiance signal better than actual local variance. To illustrate this, we inject the local variance of 2D images into AF layers without hypernetwork learning. The evaluation results in Tab. 3 show that the actual local variance is difficult to modulate the radiance signal. Also, point cloud geometry optimization is critical for some poorly reconstructed scenes. In Tab. 3, we show the quantitative results on PSNR and point scale. The preprocessing module improves performance while significantly reducing the size of the point cloud. In Fig. 8, we show the point clouds, depth maps and rendered images of Lego and Hotdog scenes. As can be seen, the noises near the crawler



	Chair	Drums	Ficus	Hotdog	Lego	Materi	Mic	Ship
PSNR								
<i>loc. var.</i>	30.25	23.48	29.90	30.33	29.49	26.55	32.22	24.47
<i>w/o. pre.</i>	32.46	25.15	30.95	33.10	27.51	28.96	33.35	26.78
ours	33.06	25.95	32.19	35.82	31.56	29.69	33.64	27.97
Number of points (Millions)								
<i>w/o. pre.</i>	0.69	0.54	0.40	1.02	1.07	1.19	0.50	1.52
ours	0.33	0.17	0.10	0.28	0.22	0.28	0.36	0.33

Table 3. Ablation analysis of adaptive frequency and point cloud geometry optimization. We inject the local variance extracted from the ground truth images as frequency information without using hypernetwork to predict frequency, denoted as *loc. var.* We also show the PSNR of training without preprocessing of point cloud geometry optimization (*w/o. pre.*) and the number of points before and after preprocessing.

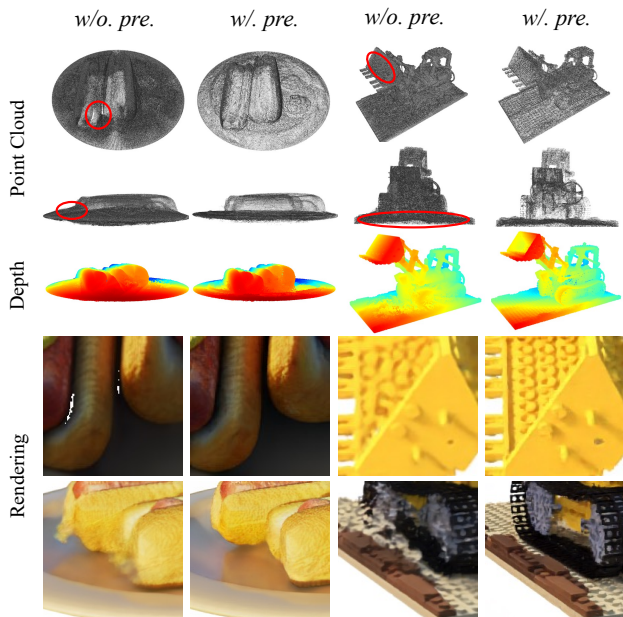


Figure 8. We show the point clouds, depth maps and rendered images with and without preprocessing of point cloud geometry optimization (*w/. pre.* and *w/o. pre.*). The noise near the crawler and the shovel in Lego scene causes artifacts and the missing part of Hotdog causes some blanks in the rendering result, while our preprocessing module optimizes the point cloud geometry and remove those artifacts.

and the shovel in Lego scene cause artifacts and the missing part of Hotdog causes some blanks in the rendering result, while our preprocessing module optimizes point cloud geometry and removes those artifacts.

### 4.3. Results of Editing

We also show the results of editing, including object editing and scene composition, which are shown in Figs. 9 and 10, respectively. As shown in Fig. 9, we edit objects in complex scenes, including rotation, translation and scaling.




Figure 9. Object-level editing on ToyDesk [55] dataset. We can scale, duplicate, translate and rotate the objects selected by users.



Figure 10. Scene composition on NeRF-Synthetic [21] and Tanks and Temples [16] datasets. These scenes only need to be trained separately. And we can edit objects while combining the scenes.

In Fig. 10, our editing algorithm handles complex occlusion relationships. Moreover, object-level editing and scene composition are performed at the same time in the last two examples in Fig. 10, *i.e.*, the balls in Materials scene are moved and composed with other scenes. More editing results can be found in supplementary material.

## 5. Conclusion

We develop a novel point cloud rendering pipeline which enables high fidelity reconstruction, real-time rendering and user-friendly editing. For the novel view synthesis task, experiments on major benchmarks demonstrate the proposed method outperforms existing point cloud rendering methods and achieves the state-of-the-art.

**Limitations and Future Work.** Our current editing method cannot relight the edited objects (*i.e.*, the transformed objects retain their original appearance), handle non-rigid object and appearance editing. Moreover, the rendering speed of the edited scenes will be reduced. We will improve these shortcomings in the future.

**Acknowledgment** This work was supported by National Science Foundation of China (U20B2072, 61976137). This work was also partly supported by SJTU Medical Engineering Cross Research Grant YG2021ZD18.



## References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *European Conference on Computer Vision*, pages 696–712. Springer, 2020. [1](#), [2](#), [3](#), [6](#), [7](#)
- [2] Nuri Benbarka, Timon Höfer, Andreas Zell, et al. Seeing implicit neural representations as fourier series. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2041–2050, 2022. [2](#)
- [3] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. *arXiv preprint arXiv:2203.09517*, 2022. [2](#), [7](#)
- [4] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. [2](#)
- [5] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. [2](#), [6](#)
- [6] Peng Dai, Yinda Zhang, Zhuwen Li, Shuaicheng Liu, and Bing Zeng. Neural point cloud rendering via multi-plane projection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7830–7839, 2020. [1](#)
- [7] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2020. [2](#), [3](#), [4](#)
- [8] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. [2](#), [7](#)
- [9] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7154–7164, 2019. [2](#)
- [10] Michelle Guo, Alireza Fathi, Jiajun Wu, and Thomas Funkhouser. Object-centric neural scene rendering. *arXiv preprint arXiv:2012.08503*, 2020. [3](#)
- [11] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016. [2](#), [4](#)
- [12] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Learning a neural 3d texture space from 2d exemplars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8356–8364, 2020. [2](#)
- [13] Xiaoyang Huang, Yi Zhang, Bingbing Ni, Teng Li, Kai Chen, and Wenjun Zhang. Boosting point clouds rendering via radiance mapping. *arXiv preprint arXiv:2210.15107*, 2022. [1](#), [2](#), [3](#), [6](#), [7](#)
- [14] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014. [2](#), [6](#)
- [15] Kacper Kania, Kwang Moo Yi, Marek Kowalski, Tomasz Trzciniński, and Andrea Tagliasacchi. Conerf: Controllable neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18623–18632, 2022. [1](#)
- [16] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017. [2](#), [6](#), [7](#), [8](#)
- [17] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. *arXiv preprint arXiv:2205.15585*, 2022. [1](#)
- [18] Zhuopeng Li, Lu Li, Zeyu Ma, Ping Zhang, Junbo Chen, and Jianke Zhu. Read: Large-scale neural scene rendering for autonomous driving. *arXiv preprint arXiv:2205.05509*, 2022. [1](#)
- [19] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020. [2](#), [7](#)
- [20] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5773–5783, 2021. [1](#)
- [21] 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. *Communications of the ACM*, 65(1):99–106, 2021. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [22] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022. [2](#), [7](#)
- [23] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015. [2](#)
- [24] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020. [2](#)
- [25] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4531–4540, 2019. [2](#)
- [26] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5589–5599, 2021. [2](#)
- [27] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation.

- In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. 2
- [28] 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*, pages 5865–5874, 2021. 2, 5
- [29] 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. *arXiv preprint arXiv:2106.13228*, 2021. 2
- [30] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14314–14323, 2021. 2
- [31] 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, 5
- [32] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pages 5301–5310. PMLR, 2019. 2, 3
- [33] Ruslan Rakhimov, Andrei-Timotei Ardelean, Victor Lempitsky, and Evgeny Burnaev. Npbg++: Accelerating neural point-based graphics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15969–15979, 2022. 1, 2, 3, 6, 7
- [34] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 6
- [35] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 1, 3
- [36] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *ACM Transactions on Graphics (TOG)*, 41(4):1–14, 2022. 1
- [37] Markus Schütz, Bernhard Kerbl, and Michael Wimmer. Rendering point clouds with compute shaders and vertex order optimization. In *Computer Graphics Forum*, volume 40, pages 115–126. Wiley Online Library, 2021. 3
- [38] Markus Schütz, Katharina Krösl, and Michael Wimmer. Real-time continuous level of detail rendering of point clouds. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 103–110. IEEE, 2019. 3
- [39] Yue Shi, Dingyi Rong, Bingbing Ni, Chang Chen, and Wenjun Zhang. Garf: Geometry-aware generalized neural radiance field. *arXiv preprint arXiv:2212.02280*, 2022. 2
- [40] Dave Shreiner, Bill The Khronos OpenGL ARB Working Group, et al. *OpenGL programming guide: the official guide to learning OpenGL, versions 3.0 and 3.1*. Pearson Education, 2009. 6
- [41] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33:7462–7473, 2020. 2, 3
- [42] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [43] Kenneth O Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162, 2007. 2
- [44] Jingxiang Sun, Xuan Wang, Yong Zhang, Xiaoyu Li, Qi Zhang, Yebin Liu, and Jue Wang. Fenerf: Face editing in neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7672–7682, 2022. 3
- [45] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2846–2855, 2021. 2
- [46] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. 2, 3
- [47] Jiaxiang Tang, Xiaokang Chen, Jingbo Wang, and Gang Zeng. Compressible-composable nerf via rank-residual decomposition. *arXiv preprint arXiv:2205.14870*, 2022. 1, 2, 3, 6, 7
- [48] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12959–12970, 2021. 2
- [49] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3835–3844, 2022. 1
- [50] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 2
- [51] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021. 2
- [52] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf:

- Point-based neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5438–5448, 2022. [2](#), [3](#), [5](#), [7](#)
- [53] Tianhan Xu and Tatsuya Harada. Deforming radiance fields with cages. In *European Conference on Computer Vision*, pages 159–175. Springer, 2022. [3](#)
- [54] Bangbang Yang, Chong Bao, Junyi Zeng, Hujun Bao, Yinda Zhang, Zhaopeng Cui, and Guofeng Zhang. Neumesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In *European Conference on Computer Vision*, pages 597–614. Springer, 2022. [3](#)
- [55] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13779–13788, 2021. [1](#), [2](#), [3](#), [7](#), [8](#)
- [56] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, pages 767–783, 2018. [2](#), [5](#)
- [57] Wang Yifan, Felice Serena, Shihao Wu, Cengiz Öztireli, and Olga Sorkine-Hornung. Differentiable surface splatting for point-based geometry processing. *ACM Transactions on Graphics (TOG)*, 38(6):1–14, 2019. [1](#)
- [58] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5752–5761, 2021. [1](#), [2](#), [7](#)
- [59] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18353–18364, 2022. [3](#)
- [60] Gizem Yüce, Guillermo Ortiz-Jiménez, Beril Besbinar, and Pascal Frossard. A structured dictionary perspective on implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19228–19238, 2022. [2](#)
- [61] Jingyang Zhang, Yao Yao, and Long Quan. Learning signed distance field for multi-view surface reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6525–6534, 2021. [2](#)
- [62] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. [5](#), [6](#)
- [63] Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Surface splatting. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 371–378, 2001. [3](#), [5](#)