

Local Implicit Ray Function for Generalizable Radiance Field Representation

Xin Huang¹, Qi Zhang², Ying Feng², Xiaoyu Li², Xuan Wang², Qing Wang¹

¹ School of Computer Science, Northwestern Polytechnical University, Xi'an 710072, China

² Tencent AI Lab

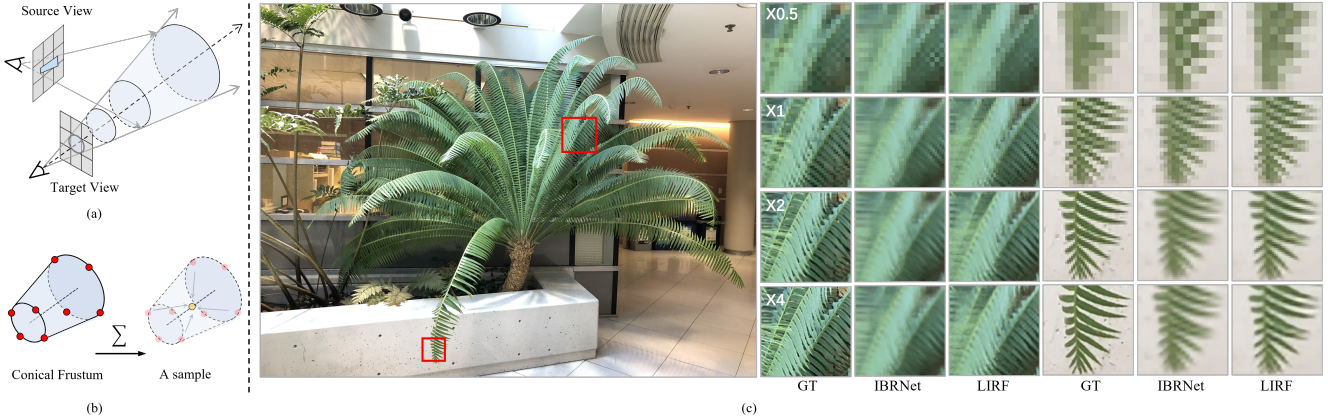


Figure 1. We propose LIRF to reconstruct radiance fields of unseen scenes for novel view synthesis. Given that current generalizable NeRF-like methods cast an infinitesimal ray to render a pixel at different scales, it causes excessive blurring and aliasing. Our method instead reasons about 3D conical frustums defined by the neighbor rays through the neighbor pixels (as shown in (a)). Our LIRF outputs the feature of any sample within the conical frustum in a continuous manner (as shown in (b)), which supports NeRF reconstruction at arbitrary scales. Compared with the previous method, our method can be generalized to represent the same unseen scene at multiple levels of details (as shown in (c)). Specifically, given a set of input views at a consistent image scale $\times 1$, LIRF enables our method to both preserve sharp details in close-up shots (anti-blurring as shown in $\times 2$ and $\times 4$ results) and correctly render the zoomed-out images (anti-aliasing as shown in $\times 0.5$ results).

Abstract

We propose **LIRF** (**Local Implicit Ray Function**), a generalizable neural rendering approach for novel view rendering. Current generalizable neural radiance fields (NeRF) methods sample a scene with a single ray per pixel and may therefore render blurred or aliased views when the input views and rendered views capture scene content with different resolutions. To solve this problem, we propose LIRF to aggregate the information from conical frustums to construct a ray. Given 3D positions within conical frustums, LIRF takes 3D coordinates and the features of conical frustums as inputs and predicts a local volumetric radiance field. Since the coordinates are continuous, LIRF renders high-quality novel views at a continuously-valued scale via volume rendering. Besides, we predict the visible weights for each input view via transformer-based feature matching to improve the performance in occluded areas. Experimental results on real-world scenes validate that our method

outperforms state-of-the-art methods on novel view rendering of unseen scenes at arbitrary scales.

1. Introduction

Novel view synthesis has garnered recent attention with compelling applications of neural rendering in virtual and augmented reality. Different from image-based rendering [6, 19, 32, 42, 74], Neural Radiance Fields (NeRF) [43] implicitly represents the 3D scenes within multilayer perceptrons (MLPs) by mapping coordinates to color and geometry of scenes. To render a pixel, the ray projected to that pixel is traced and the color of each sampled point along the ray is accumulated based on volume rendering.

Despite NeRF and its variants having demonstrated remarkable performance in providing immersive experiences in various view synthesis tasks, their practical applications are constrained by the requirement of training from scratch on each new scene, which is time-consuming. To overcome

this problem, many researches [10, 12, 29, 33, 37, 58, 64, 70] introduce image-based rendering techniques to NeRF, which achieves generalization on unseen scenes. They take into consideration the image features (from nearby views) of a 3D point. The common motivation is to predict the density and color of this point by matching the multi-view features, which is similar to the stereo matching methods [20, 54, 68] that find a surface point by checking the consistency of multi-view features.

While these methods generalize well on new scenes when the distance of input and testing views are roughly constant from the scene (as in NeRF), they cannot properly deal with the less constrained settings such as different resolutions or varying focal length and produce results with blurring or aliasing artifacts. Since a single ray is cast through each pixel whose size and shape are ignored, it's challenging to query the accurate feature of the target ray from input images as shown in Fig. 2(a), and the model learns an ambiguous result as shown in Fig. 2(b). Mip-NeRF [3], a NeRF variant of per-scene optimization, proposes an anti-aliasing design that models the ray through a pixel as a cone and uses a 3D Gaussian to approximate the sampled *conical frustum* (a cone cut perpendicular to its axis) for volumetric representation. However, directly extending Mip-NeRF to a generalizable method is also challenging to extract the accurate features of the ray from input images due to the subpixel precision. Consequently, an efficient solution is to supersample each pixel by marching multiple rays according to its footprint, similar to the strategy used in offline raytracing.

Our key insight is the *local implicit ray function* (LIRF) that represents the feature aggregation of samples within ray conical frustum in a continuous manner, as shown in Fig. 1. Specifically, given any 3D sampled position within a conical frustum, our LIRF outputs the aggregated feature by taking the 3D coordinate of this position and the features of vertices within the conical frustum as inputs (the vertices of a conical frustum are defined with eight points (red points) as shown in Fig. 1). The continuous sampled position allows our method to arbitrarily upsample the rendered rays and thus synthesize novel views of the same unseen scene at multiple levels of detail (anti-blurring and anti-aliasing). Furthermore, recent generalizable NeRF methods [29, 37] introduce multi-view depth estimation to reduce the artifacts caused by occlusions, but it is computationally expensive to construct the cost volume for each view. We instead match local multi-view feature patches to estimate the visibility weights of each sample for anti-occlusion. Overall, our main contributions are:

1. A new generalizable approach that renders pixels by casting cones and outperforms existing methods on novel view synthesis at multiple scales.
2. A local implicit ray function that simplifies the repre-




Figure 2. Most generalizable variants of NeRF represent a ray as a set of infinitesimal samples (shown here as dots) along that ray and map these samples into input views to query image features for volumetric representation prediction. However, this results in two drawbacks when training on multi-scale images with less constrained settings: (a) Inaccurate features. The sampling strategy which ignores the shape and size of each ray is difficult to query accurate image features. (b) Ambiguous supervisions. The same 3D position captured by cameras under different scales results in different colors because these pixels are the integral of regions with different shapes and sizes (shown here as trapezoids). During the training, the network learns to map the same image features (from the source view) to these different colors, which causes ambiguous results.

sentation of conical frustums and enables continuous supersampling of rays.

3. A transformer-based visibility weight estimation module that alleviates the occlusion problem.

To evaluate our method, we construct extensive series of experiments on real forward-facing scenes. Our experiments show that trained on large amounts of multi-view data, LIRF outperforms state-of-the-art generalizable NeRF-like methods on novel views synthesis for unseen scenes.

2. Related Work

Image-based rendering. When the images for a scene are captured densely, earlier lines of work directly interpolate novel views from input views [19, 32] or render views by weighted blending [6, 13]. More recently, researchers have been focused on the issue of novel views synthesis from sparse inputs. To map and blend the input images in a novel target view, some methods [7, 44] use the proxy geometry obtained via structure-from-motion (SfM) or multi-view stereo (MVS). Moreover, to improve the accuracy of mapping and blending, some improved methods are proposed for proxy geometry estimation [9, 24], optical flow correction [15, 16] and deep blending [23, 49, 50]. Differently, some other methods directly reconstruct the texture of mesh [14, 25, 60] or points cloud [1, 41, 47, 52] for novel view synthesis. However, the rendering quality by IBR methods is directly affected by geometric inaccuracies of 3D reconstruction methods [27, 53], especially in low-textured or reflective regions.

Explicit volumetric representations. The other methods render high-quality novel views by constructing an explicit scene representation, such as voxel grids [30, 62] and layered-depth images [46, 55], from captured images. For example, Zhou *et al* [74] represent a scene with a set of depth-dependent images with transparency, which is called multi-plane images (MPIs). Novel views can be directly rendered from MPIs using alpha compositing. Recently, MPI representation has been applied in many scenarios, such as rendering novel views from multiple MPIs [42], novel view synthesis from a single image [34, 61], variants as multi-sphere images [2, 5], and 3D image generation [67, 73]. Trained on a large multi-view dataset, MPI-based methods can generalize well to unseen scenes and render photo-realistic novel views fast. However, those methods always render views within limited viewing volumes. Besides, the volumetric representations explicitly decompose a scene into extensive samples, which requires large memory to store them and limits the resolution of novel views.

Neural scene representations. Representing the geometry and appearance of a scene with neural networks has been a surge. Traditional methods explicitly represent scenes with point clouds [52, 65], meshes [25, 60], or voxels [38, 56], while neural scene representations implicitly encode the scene with continuous coordinate-based functions such as radiance fields [3, 43], signed distance fields [8, 18, 69], or occupancy fields [40, 45]. NeRF [43] approximates a 3D scene with an implicit function and has shown high-quality view synthesis results. Mip-NeRF [3] reduces objectionable aliasing artifacts and improves NeRF’s ability to represent scene details by casting a cone instead of a ray. While NeRF has been expanded to many new scenarios [4, 11, 21, 26, 35, 39, 48, 57, 66, 71], most NeRF-like models take coordinate as inputs, which restricts their ability to generalize to unseen scenes. Recently, some generalizable NeRF-like methods have been proposed [10, 12, 29, 33, 37, 64, 70]. PixelNeRF [70], SRP [12], MINE [33] and MVSNerF [10] try to construct radiance fields on-the-fly from sparse input views, while they struggle with complex scenes. To solve occlusions, NeuRay [37] and GeoNeRF [29] estimate depth using MVS methods [20, 68] and calculate the occlusion masks from the estimated depth maps. However, it’s expensive to construct cost volumes per source view to estimate depth. While these generalizable methods have been capable of rendering novel views for unseen scenes, rendering novel views by casting a single ray may produce renderings that are aliased or blurred.

3. Method

Our goal is to predict volumetric radiance fields from a set of multi-view images captured at a consistent image scale ($\times 1$), and output novel views at continuous scales ($\times 0.5 \sim \times 4$). Our proposed framework is presented in

Fig. 3, which is composed of five parts: 1) extracting 2D feature maps from source images (Sec. 3.2), 2) obtaining the image feature for the samples on target rays via local implicit ray function (Sec. 3.3), 3) predicting the visibility weights of each source view by matching feature patches (Sec. 3.4), 4) aggregating local image features from different source views and mapping them into colors and densities (Sec. 3.5), 5) rendering a target pixel via volume rendering (Sec. 3.1).

3.1. Volume Rendering

We first briefly review the volume rendering in NeRF [43]. NeRF implicitly encodes a scene as a continuous radiance field. Given a ray \mathbf{r} passing through the scene, NeRF maps the 5D coordinates (3D for position, 2D for direction) of each sampled location into view-dependent color and volume density via an MLP. To determine the pixel color of the ray \mathbf{r} , NeRF first obtains the colors and densities of all samples on the ray \mathbf{r} and then accumulates the colors according to densities. The volume rendering thus is defined by:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad (1)$$

$$T_i = \exp\left(-\sum_{l=1}^{i-1} \sigma_l \delta_l\right),$$

where N is the number of samples along a ray with ascending depth values. The \mathbf{c} and σ denote colors and densities, respectively. T denotes the accumulated transmittance. The δ is the distance between adjacent samples. Following IBRNet [64], we remove δ in volume rendering for a better generalization. To render a novel view of unseen scenes, we will introduce our proposed method to obtain colors and densities from input images instead of coordinates.

3.2. Image Feature Extraction

Given a target viewpoint, we select a set of V nearby source images and their corresponding camera poses as inputs. Our method relies on the local image features from source images to produce target images. Usually, most generalizable methods [37, 64] extract dense image features from each source view by a U-Net [51] architecture with ResNet [22]. Although U-Net and ResNet are widely adopted in high-level computer vision tasks such as image recognition and image segmentation, our rendering is based on pixel-wise operations. The feature extraction network is supposed to pay more attention to image details, so we use an EDSR network [36] to extract image features. The EDSR network removes unnecessary modules in conventional residual networks, which makes the network more focused on detailed content. We modify the last layer in the




Figure 3. The overview of LIRF. To render a target image, a set of V neighboring source views are selected and the 2D feature map of each source view is extracted via an EDSR network [36]. N points (yellow points) are sampled along a target ray, and the corresponding conical frustum is represented with M vertices (red points). (a) For a sample on target ray at 3D position \mathbf{x} , we obtain its feature \mathbf{z} by aggregating the features of the vertices and two latent codes (relative position $\Delta\mathbf{x} = \mathbf{x} - \mathbf{q}$ and target scale s). (b) To deal with occlusions, we also estimate the visibility weights of each source view. Visibility weights can be considered the probability of a sample being observed by each source view. Taking the image features \mathbf{f}^w , source ray direction \mathbf{d} and target scale s as inputs, our model outputs the visibility weights in source views of sample \mathbf{x} . (c) After obtaining the features \mathbf{z} and corresponding visibility weights w , we aggregate them by a transformer [63] module \mathcal{T}_2 which outputs two parts of features: one for the colors and the other for densities. Next, we explicitly blend features \mathbf{t}_1' and map them into colors \mathbf{c} by an MLP \mathcal{M}_c . As for features \mathbf{t}_2' , an AE network [29] aggregates their information along a ray and an MLP \mathcal{M}_σ maps the features to densities σ . (d) Finally, the color of target ray is rendered by volume rendering from the colors \mathbf{c} and densities σ . The mean squared error between the predicted color and ground truth color is calculated for optimization.

EDSR network to output two image features for the following modules, with one image feature for the estimation of visibility weights, and the other for the prediction of colors and densities. Given a source image I , the image feature extraction is formulated as:

$$(\mathbf{F}^w, \mathbf{F}^c) = \text{EDSR}(I), \quad (2)$$

where \mathbf{F}^w denotes the feature map used for visibility weights and \mathbf{F}^c denotes the feature map used for colors and densities.

3.3. Local Implicit Ray Function

In our method, a target ray is defined as the integral of a conical volume. To consider the shape and size of the volume viewed by each ray, we define the feature of each sample along the ray by a continuous integration within the corresponding conical frustum. Given the 3D position \mathbf{x} within the conical frustum, its feature \mathbf{z} is defined as:

$$\mathbf{z} = \int W(\|\Delta\mathbf{x}_q\|_2) \mathcal{R}(\mathbf{f}_q, \Delta\mathbf{x}_q) d\mathbf{q}, \quad (3)$$

where \mathbf{q} denotes a 3D position within the conical frustum, \mathbf{f}_q is the feature at position \mathbf{q} , $\Delta\mathbf{x}_q = \mathbf{x} - \mathbf{q}$ denotes the

relative position between two points, and $\|\cdot\|_2$ denotes the Euclidean distance. W is a weights function that outputs weights according to the distance $\|\Delta\mathbf{x}_q\|_2$. \mathcal{R} is an MLP-based function for the aggregation of features and positions. **Discrete representation.** Considering the computation and memory consumption, we cast four rays from the cone's footprint to discretely approximate the continuous conical volume. Inspired by Plenoxels [17], a voxel-based representation, each conical frustum along the cone is represented with 8 vertices, and the target ray passes through all conical frustums, as shown in Fig. 4(b). Any samples within the conical frustum can be obtained by trilinear interpolation, which is similar to the sampling strategy of voxel-based representation. Specifically, we project the vertices on the four rays into the source views and query corresponding features \mathbf{f}^c from feature maps \mathbf{F}^c . For a sample at position \mathbf{x} within the corresponding conical frustum, its feature is defined as:

$$\mathbf{z} = \sum_{j=1}^M W(\|\Delta\mathbf{x}_j\|_2) \mathcal{R}(\mathbf{f}_j^c, \Delta\mathbf{x}_j), \quad (4)$$

$$W(\|\Delta\mathbf{x}_j\|_2) = \frac{\|\Delta\mathbf{x}_j\|_2}{\sum_{l=1}^M \|\Delta\mathbf{x}_l\|_2},$$

where $M = 8$ is the number of vertices used to represent a conical frustum. \mathbf{f}^c is the image feature of the vertices on the conical frustum. $\Delta\mathbf{x}$ denotes the relative position between the sample and the vertices within the conical frustum.

Cone radius. A cone cast from the view origin and passing through the image plane, the radius of the cone at the image plane is called cone radius and parameterized as r . The cone radius is affected by image resolutions and observation distances (being closer to objects means a larger scale). However, it is hard to know the observation distances beforehand. Besides, since our testing scenes are captured using a hand-held camera, the observation distance of novel views and source views are slightly different. To adjust the cone size accurately and conveniently, we don't directly modify the cone size only according to the relative image resolution. Instead, we set a maximum cone radius r_m which is the pixel width of the target image with the minimum resolution, and a latent code is introduced to implicitly control the size of a cone. Specifically, we first cast four rays passing through the four red points in the target image plane, as shown in Fig. 4(a). For any target ray projecting within the circle with radius r_m , its features are aggregated from the features of the four neighboring rays. Moreover, a scale-dependent input \mathbf{s} is additionally introduced to modify the cone size, where \mathbf{s} is the relative scale of a target image. We reformulate Eq. (4) with the form:

$$\mathbf{z} = \sum_{j=1}^M W(\|\Delta\mathbf{x}_j\|_2) \mathcal{R}(\mathbf{f}_j^c, \Delta\mathbf{x}_j, \mathbf{s}). \quad (5)$$

In practice, the positional encoding strategy [43] is applied on the relative position $\Delta\mathbf{x}$ but not on relative scale \mathbf{s} , since we want the implicit function to learn more high-frequency information of image feature and adjust the scale of renderings smoothly.

3.4. Visibility Weights Estimation

In this stage, we estimate the visibility weight that reveals the occlusions in each source view. Without any geometric priors, it's challenging to solve the occlusions thoroughly. As a compromise, we assume that the occluded content appears in most source views. Our model estimates the visibility weights of one source view by matching its feature patch with the feature patches from other source views. For the samples on a target ray, we obtain their feature patches \mathbf{f}^w from the image features \mathbf{F}^w as shown in Fig. 4(c). Apart from the feature patches, two latent codes are introduced to improve the performance. To consider the direction for each source ray, their directions \mathbf{d} in target camera coordinates are input. We do not use the global directions of source rays in world coordinates, since different scenes have different world coordinates and local directions are more suitable for the generalization to unseen scenes.




Figure 4. Method details of LIRF. The \mathbf{x} denotes the position of a sample, \mathbf{q} denotes the position of a vertex, and \mathbf{d} denotes the direction of a source ray. (a) The rays projecting at target image plane. The size of target image is normalized. (b) A cone is represented with four rays (solid lines, only two are drawn) and a conical frustum is represented with eight vertices (red points). A target ray passes through the cone (dotted line). (c) The feature patches for visibility weights are obtained by warping samples on the target ray to feature maps.

To consider the scale of the target image, we also introduce the target scale \mathbf{s} . We define the feature concatenation of all inputs as:

$$\mathbf{t}_0 = \text{MLP}(\mathbf{f}^w \parallel \mathbf{d} \parallel \mathbf{s}), \quad (6)$$

where \parallel denotes the concatenation operation and MLP denotes a two-layer MLP used to reduce the channels of features. Positional encoding [43] is applied on directions \mathbf{d} .

Next, a self-attention transformer [63] module \mathcal{T}_1 is used to fully match these aggregated feature \mathbf{t}_0 and output new features for visibility weights. Formally, this stage is written as:

$$\mathbf{t}_0' = \mathcal{T}_1(\mathbf{t}_0). \quad (7)$$

We then map the new feature \mathbf{t}_0' into visibility weights w by an MLP \mathcal{M}_w ,

$$w = \mathcal{M}_w(\mathbf{t}_0'). \quad (8)$$

3.5. Color and Density Prediction

After obtaining the local image features \mathbf{z} and the visibility weights w , the colors and densities can be predicted. We aggregate the features and visibility weights to features \mathbf{t}_1 . Besides, the weighted mean and weighted variance of features \mathbf{z} over all source views are calculated based on visibility weights w . The mean and variance are then aggregated into features \mathbf{t}_2 . This process is defined as:

$$\begin{aligned} \mathbf{t}_1 &= \text{MLP}(\mathbf{z} \parallel w), \\ \mathbf{t}_2 &= \text{MLP}(\text{mean}(\mathbf{z}, w) \parallel \text{var}(\mathbf{z}, w)). \end{aligned} \quad (9)$$

\mathbf{t}_1 could be considered as the color information from source views, while \mathbf{t}_2 could be considered as the measure of image feature consistency. The densities are predicted by checking the feature consistency, since local image features from different source views are usually consistent when the sample is on the object surface [37].

Next, the two features are fed into a self-attention transformer \mathcal{T}_2 for fully information aggregation,

$$(\mathbf{t}_1', \mathbf{t}_2') = \mathcal{T}_2(\mathbf{t}_1, \mathbf{t}_2). \quad (10)$$

The feature \mathbf{t}_1' and \mathbf{t}_2' now have combined the information from the the local image features and visibility weights. Similar to most generalizable NeRF-like methods [29, 37, 64], the density is estimated by an auto-encoder-style network [29] AE which aggregates information along a ray, and an MLP \mathcal{M}_σ that maps features to densities,

$$\sigma = \mathcal{M}_\sigma(\text{AE}(\mathbf{t}_2')). \quad (11)$$

As for colors, we further explicitly blend the features \mathbf{t}_1' according to the visibility weights, and then an MLP \mathcal{M}_c are used for color prediction,

$$\mathbf{c} = \mathcal{M}_c\left(\sum_{k=1}^V (\mathbf{t}_{1',k} w_k)\right). \quad (12)$$

Once the densities and colors of the samples on target rays are predicted, the final colors of target rays are rendered via volume rendering (Eq. (1)).

3.6. Loss Function

We minimize the mean squared error between our predicted colors and ground truth colors for optimization:

$$\mathcal{L} = \sum_{\mathbf{r} \in \Omega} \|\hat{C}(\mathbf{r}) - C(\mathbf{r})\|_2^2, \quad (13)$$

where Ω is a set of camera rays at target position. \hat{C} and C are the predicted colors and ground truth colors, respectively.

4. Experiments

4.1. Implementation Details

To render a target view, we select nearby source views using $V = 8$ during both training and testing. $M = 8$ vertices are used to represent a conical frustum. The cone radius r_m is set to the pixel width of the target image with the minimum resolution. We directly sample $N = 128$ points on a ray instead of using the coarse-to-fine sampling [43]. The extracted image feature maps have 32 channels and their size is the same as the original image size. The size of feature patches used for visibility weights is set to 7×7 . We train the generalizable LIRF for 250k iterations. The training batch size of rays is set to 512. We use Adam optimizer [31] with a learning rate of 5×10^{-4} that decays exponentially along with the optimization. We train our model on eight V100 GPUs, which takes about two days.

4.2. Datasets

Our model is trained on three real datasets: real DTU multi-view dataset [28] and two real forward-facing datasets from LLFF [42] and IBRNet [64]. All 190 scenes (35 scenes from LLFF, 67 scenes from IBRNet and 88 scenes from DTU dataset) are used for training. Eight unseen scenes from the LLFF dataset are used for testing. During the multi-scale training, the resolutions of all input views are consistent, while the resolution of each target view is randomly selected from 1 to 4 times the input resolution. When training our model on single-scale datasets, the image resolutions of input and target images are the same. During testing, the resolution of input views is 504×378 . We evaluate our model by rendering novel views at multiple scales: $\times 0.5$, $\times 1$, $\times 2$ and $\times 4$ ($\times 0.5$ denotes 0.5 times the resolution of input views, and so on).

4.3. Results

We evaluate our model on LLFF testing scenes and compare it against three state-of-the-art generalizable methods: IBRNet [64], NeuRay [37] and GeoNeRF [29]. For quantitative evaluations, we report three error metrics, PSNR, SSIM and LPIPS [72]. Following prior work [3, 59], we also summarize all three error metrics into an ‘‘average’’ error metric by calculating the geometric mean of $\text{MSE} = 10^{-\text{PSNR}/10}$, $\sqrt{1 - \text{SSIM}}$, and LPIPS, which provides an easier comparison.

Multi-scale novel views. The performance of our method on rendering multi-scale novel views can be seen in Tab. 1 and Fig. 5. As shown in Tab. 1, our method outperforms baseline methods on all four scales. Though our model isn’t trained on $\times 0.5$ scale, it also can render low-scale views by explicitly modifying the cone radius r_m according to the pixel size at $\times 0.5$ scale. In contrast, the baselines fail to render view at a lower scale due to the aliasing artifacts (shown as the $\times 0.5$ results by IBRNet in Fig. 1(c)). We notice that the baselines produce relatively better results on rendering novel views at an up-scale after training on multi-scale datasets. On the other hand, they produce worse results on low-scale ($\times 0.5$ and $\times 1$) compared with their released models, since baselines have difficulty converging well when supervisions have different image scales, as discussed in Fig. 2 (b). Besides, our dataset contains fewer training scenes (about 15% of the training scenes used by IBRNet or NeuRay). The qualitative comparisons are shown in Fig. 5. It can be seen that the improvements produced by our method are most visually apparent in challenging cases such as small or thin structures. Besides, our renderings have fewer blurring artifacts and are closer to the ground truth.

Single-scale novel views. We also train our model on the single-scale dataset to evaluate our ability on novel view synthesis. The quantitative results are shown in Tab. 2. One

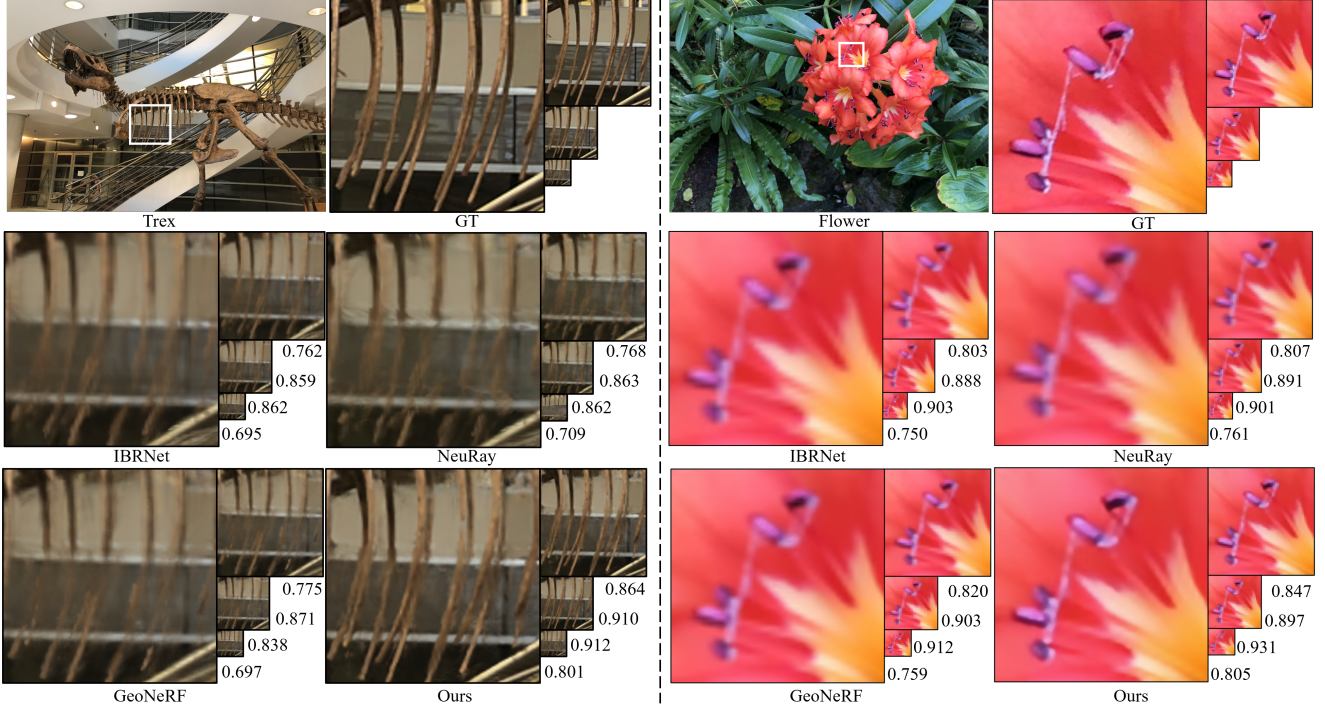


Figure 5. Qualitative comparisons on real forward-facing scenes. The cropped regions of two scenes are rendered at different scales ($\times 0.5$, $\times 1$, $\times 2$ and $\times 4$). SSIM averaged over all novel views for each scale are provided at the lower right of images.

Table 1. Quantitative comparisons of LIRF and its ablations against IBRNet [64], NeuRay [37] and GeoNeRF [29] on LLFF multi-scale testing dataset. Metrics are averaged over all scenes. * denotes training on the same multi-scale training set as our method.

	PSNR \uparrow				SSIM \uparrow				LPIPS \downarrow				Avg. \downarrow
	$\times 0.5$	$\times 1$	$\times 2$	$\times 4$	$\times 0.5$	$\times 1$	$\times 2$	$\times 4$	$\times 0.5$	$\times 1$	$\times 2$	$\times 4$	
IBRNet	25.06	25.28	23.18	22.15	0.866	0.840	0.731	0.669	0.108	0.160	0.299	0.442	0.079
NeuRay	24.80	25.17	23.10	22.08	0.859	0.837	0.729	0.667	0.107	0.157	0.294	0.434	0.080
GeoNeRF	24.89	25.74	23.43	22.27	0.864	0.864	0.753	0.679	0.108	0.136	0.274	0.421	0.076
IBRNet*	22.96	23.62	22.33	21.51	0.816	0.813	0.723	0.665	0.140	0.178	0.307	0.444	0.090
NeuRay*	22.79	22.40	21.23	20.61	0.794	0.733	0.646	0.622	0.172	0.262	0.382	0.493	0.107
GeoNeRF*	23.39	25.08	23.81	22.69	0.821	0.859	0.784	0.708	0.138	0.134	0.255	0.401	0.077
Ours	26.75	25.93	24.58	23.79	0.905	0.877	0.816	0.760	0.100	0.124	0.227	0.373	0.063
Ours (single ray)	25.38	25.72	24.27	23.43	0.881	0.871	0.794	0.723	0.135	0.132	0.254	0.392	0.069
Ours w/o scale	26.81	25.90	24.52	23.72	0.904	0.873	0.807	0.749	0.098	0.128	0.235	0.374	0.064
Ours w/o position	25.95	25.72	24.36	23.56	0.882	0.869	0.803	0.741	0.127	0.136	0.242	0.373	0.067
Ours w/o patch	26.85	25.95	24.56	23.76	0.906	0.877	0.813	0.756	0.097	0.127	0.231	0.367	0.064
Ours w/o direction	26.35	25.36	24.10	23.38	0.899	0.864	0.799	0.744	0.103	0.137	0.242	0.373	0.067
Ours w/o vis. weights	25.90	25.11	23.91	23.19	0.888	0.856	0.792	0.737	0.118	0.148	0.250	0.378	0.070
Ours (U-Net feat.)	26.00	25.10	23.43	22.77	0.887	0.853	0.761	0.706	0.116	0.151	0.298	0.427	0.075

can see that training on the single-scale dataset does not reduce our performance on rendering novel views at $\times 1$ scale, and even improves our results on PSNR and SSIM metrics. Figure 6 presents the visualized comparisons on cropped regions. Although training on the single-scale dataset, our method shows a visual improvement on the challenging objects, such as the thin leaves.

Visibility weights. To demonstrate the effectiveness of our

method in dealing with occlusions, we visualize the estimated visibility weights. As shown in Fig. 7, region A is occluded in source view 1, so it has smaller visibility weights in region A. Conversely, region A is visible in source view 2, while region B is occluded. Therefore, the visibility weights of source view 2 have larger values at region A and smaller values at region B. With visibility weights estimation, our method correctly renders the scene content at




Figure 6. Qualitative comparisons on single-scale novel view synthesis. Though trained on the single-scale dataset, Ours’ still outperforms baselines. SSIM averaged over all novel views are provided at lower right.

Table 2. Quantitative comparisons of LIRF trained on single-scale dataset against baseline methods on rendering novel view at $\times 1$ scale. “Ours’” denotes our method training on single-scale dataset, while “Ours” denotes our method training on multi-scale dataset.

	IBRNet	NeuRay	GeoNeRF	Ours’	Ours
PSNR \uparrow	25.28	25.17	25.74	26.30	25.93
SSIM \uparrow	0.840	0.837	0.864	0.878	0.877
LPIPS \downarrow	0.160	0.157	0.136	0.128	0.124

regions A and B.

4.4. Ablation Studies

Ablation studies are shown in Tab. 1 to investigate the individual contribution of key modules of our model. (a) For “single ray”, we modify our model to render a pixel from a single ray instead of a cone and keep everything else the same. Compared with our full model, the performance by “single ray” is reduced on all four testing scales, especially on the $\times 0.5$ scale and $\times 4$ scale, which demonstrates the contribution of our local implicit ray function for rendering multi-scale novel views. (b) For “w/o scale”, we remove the scale s in Eq. (5) and Eq. (6). It reduces our model’s ability to modify our implicit ray function according to the scale of target views. (c) For “w/o position”, the relative position Δx in Eq. (5) is removed, which prevents our implicit ray function from perceiving the relative 3D position between a sample and the vertices of conical frustum. (d) For “w/o patch”, the size of the feature patches used to predict visibility weights is set to 1×1 . This reduces our performance on rendering novel views with higher scales while improving our model on rendering $\times 0.5$ novel views. When testing on different scales, the patch size is set to 7×7 , which may be too large to render views at $\times 0.5$ scale. (e) For “w/o direction”, we remove the direction d in Eq. (6). Without considering the direction of source rays, our model produces worse results on all testing scales. (f) For “w/o vis. weights”, we remove the visibility weights estimation module. Without visibility weights, our model fails to solve the occlusions, which significantly reduces our




Figure 7. Visualizations of visibility weights (only two source views are presented). Our model accurately estimates the invisible regions for source views. For example, region A is occluded in source view 1, thus the visibility weights at this region are “dark”.

performance. (g) For “U-Net feat.”, the image feature maps are extracted by the U-Net of IBRNet [64]. We can see that our performance is reduced a lot. Most recent generalizable methods focus on the framework design of predicting color and densities from image features. Actually, the image feature extraction network is also important, affecting the quality of novel views from the source.

5. Conclusion

We have proposed a novel method for novel view synthesis of unseen scenes. Our method not only renders novel views with fewer blurring artifacts, but also produces novel views at arbitrary scales, even at higher scales than input views. The core of our method is the local implicit ray function that constructs each target ray by aggregating the information of a cone. Besides, our method also estimates the visibility weights to mitigate the occlusion problem. Compared with prior works, our method performs better on novel view synthesis of unseen scenes. Our code and models will be released to the research community to facilitate reproducible research.

Limitations. Rendering novel views of unseen scenes is a challenging task. Similar to most image-based rendering methods, our method renders novel views from selected local input views, which causes two problems. First, our method can’t render the regions occluded in all input views. Second, since selected input views vary with the novel view position, compared with NeRF, our method is weak in view consistency. Moreover, our method renders a pixel by casting four rays to approximate a cone, which increases the computational cost. Finally, though our method can mitigate the artifacts caused by occlusions, it fails in some challenging scenes. To solve this, a possible choice is introducing geometric priors.

Acknowledgements. The work was supported by NSFC under Grant 62031023.

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *ECCV*, pages 696–712, 2020. [2](#)
- [2] Benjamin Attal, Selena Ling, Aaron Gokaslan, Christian Richardt, and James Tompkin. Matryodshka: Real-time 6dof video view synthesis using multi-sphere images. In *ECCV*, 2020. [3](#)
- [3] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, pages 5855–5864, 2021. [2](#), [3](#), [6](#)
- [4] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. NeRD: Neural reflectance decomposition from image collections. In *ICCV*, pages 12684–12694, 2021. [3](#)
- [5] Michael Broxton, John Flynn, Ryan Overbeck, Daniel Erickson, Peter Hedman, Matthew Duvall, Jason Dourgarian, Jay Busch, Matt Whalen, and Paul Debevec. Immersive light field video with a layered mesh representation. *ACM TOG*, 39(4):86–1, 2020. [3](#)
- [6] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *SIGGRAPH*, pages 425–432, 2001. [1](#), [2](#)
- [7] Rodrigo Ortiz Cayon, Abdelaziz Djelouah, and George Drettakis. A bayesian approach for selective image-based rendering using superpixels. In *3DV*, pages 469–477, 2015. [2](#)
- [8] Rohan Chabra, Jan E Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *ECCV*, pages 608–625, 2020. [3](#)
- [9] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM TOG*, 32(3):1–12, 2013. [2](#)
- [10] 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 *ICCV*, pages 14124–14133, 2021. [2](#), [3](#)
- [11] Xingyu Chen, Qi Zhang, Xiaoyu Li, Yue Chen, Ying Feng, Xuan Wang, and Jue Wang. Hallucinated neural radiance fields in the wild. In *CVPR*, pages 12943–12952, 2022. [3](#)
- [12] Julian Chibane, Aayush Bansal, Verica Lazova, and Gerard Pons-Moll. Stereo radiance fields (srf): Learning view synthesis for sparse views of novel scenes. In *CVPR*, pages 7911–7920, 2021. [2](#), [3](#)
- [13] Abe Davis, Marc Levoy, and Fredo Durand. Unstructured light fields. In *Comput. Graph. Forum*, volume 31, pages 305–314, 2012. [2](#)
- [14] Paul Debevec, Yizhou Yu, and George Borshukov. Efficient view-dependent image-based rendering with projective texture-mapping. In *Eurographics Workshop on Rendering Techniques*, pages 105–116. Springer, 1998. [2](#)
- [15] Ruofei Du, Ming Chuang, Wayne Chang, Hugues Hoppe, and Amitabh Varshney. Montage4d: interactive seamless fusion of multiview video textures. In *13D*, 2018. [2](#)
- [16] Martin Eisemann, Bert De Decker, Marcus Magnor, Philippe Bekaert, Edilson De Aguiar, Naveed Ahmed, Christian Theobalt, and Anita Sellent. Floating textures. *Comput. Graph. Forum*, 27(2):409–418, 2008. [2](#)
- [17] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, pages 5501–5510, 2022. [4](#)
- [18] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3d shape. In *CVPR*, pages 4857–4866, 2020. [3](#)
- [19] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. In *SIGGRAPH*, pages 43–54, 1996. [1](#), [2](#)
- [20] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *CVPR*, pages 2495–2504, 2020. [2](#), [3](#)
- [21] Michelle Guo, Alireza Fathi, Jiajun Wu, and Thomas Funkhouser. Object-centric neural scene rendering. *arXiv preprint arXiv:2012.08503*, 2020. [3](#)
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [3](#)
- [23] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM TOG*, 37(6):1–15, 2018. [2](#)
- [24] Peter Hedman, Tobias Ritschel, George Drettakis, and Gabriel Brostow. Scalable inside-out image-based rendering. *ACM TOG*, 35(6):1–11, 2016. [2](#)
- [25] Jingwei Huang, Justus Thies, Angela Dai, Abhijit Kundu, Chiyu Jiang, Leonidas J Guibas, Matthias Nießner, Thomas Funkhouser, et al. Adversarial texture optimization from rgb-d scans. In *CVPR*, pages 1559–1568, 2020. [2](#), [3](#)
- [26] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. HDR-NeRF: High dynamic range neural radiance fields. In *CVPR*, pages 18398–18408, 2022. [3](#)
- [27] Michal Jancosek and Tomáš Pajdla. Multi-view reconstruction preserving weakly-supported surfaces. In *CVPR*, pages 3121–3128, 2011. [2](#)
- [28] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engil Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *CVPR*, pages 406–413, 2014. [6](#)
- [29] Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. GeoNeRF: Generalizing nerf with geometry priors. In *CVPR*, pages 18365–18375, 2022. [2](#), [3](#), [4](#), [6](#), [7](#)
- [30] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. *NeurIPS*, 30, 2017. [3](#)
- [31] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. [6](#)
- [32] Marc Levoy and Pat Hanrahan. Light field rendering. In *SIGGRAPH*, pages 31–42, 1996. [1](#), [2](#)
- [33] Jiaxin Li, Zijian Feng, Qi She, Henghui Ding, Changhu Wang, and Gim Hee Lee. Mine: Towards continuous depth mpi with nerf for novel view synthesis. In *ICCV*, pages 12578–12588, 2021. [2](#), [3](#)

- [34] Qinbo Li and Nima Khademi Kalantari. Synthesizing light field from a single image with variable mpi and two network fusion. *ACM TOG*, 39(6), 2020. 3
- [35] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, pages 6498–6508, 2021. 3
- [36] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPRW*, pages 136–144, 2017. 3, 4
- [37] Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou, and Wenping Wang. Neural rays for occlusion-aware image-based rendering. In *CVPR*, pages 7824–7833, 2022. 2, 3, 5, 6, 7
- [38] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM TOG*, 38(4):65:1–65:14, 2019. 3
- [39] Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-NeRF: Neural radiance fields from blurry images. In *CVPR*, pages 12861–12870, 2022. 3
- [40] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, pages 4460–4470, 2019. 3
- [41] Moustafa Meshry, Dan B Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural rendering in the wild. In *CVPR*, pages 6878–6887, 2019. 2
- [42] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM TOG*, 38(4):1–14, 2019. 1, 3, 6
- [43] 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 *ECCV*, pages 405–421, 2020. 1, 3, 5, 6
- [44] Ryan S Overbeck, Daniel Erickson, Daniel Evangelakos, Matt Pharr, and Paul Debevec. A system for acquiring, processing, and rendering panoramic light field stills for virtual reality. *ACM TOG*, 37(6):1–15, 2018. 2
- [45] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *ECCV*, pages 523–540, 2020. 3
- [46] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM TOG*, 36(6):1–11, 2017. 3
- [47] Francesco Pittaluga, Sanjeev J Koppal, Sing Bing Kang, and Sudipta N Sinha. Revealing scenes by inverting structure from motion reconstructions. In *CVPR*, pages 145–154, 2019. 2
- [48] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. In *CVPR*, pages 10318–10327, 2021. 3
- [49] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *ECCV*, pages 623–640, 2020. 2
- [50] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *CVPR*, pages 12216–12225, 2021. 2
- [51] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 3
- [52] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *ACM TOG*, 41(4):1–14, 2022. 2, 3
- [53] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, pages 4104–4113, 2016. 2
- [54] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, pages 501–518, 2016. 2
- [55] Jonathan Shade, Steven Gortler, Li-wei He, and Richard Szeliski. Layered depth images. In *SIGGRAPH*, pages 231–242, 1998. 3
- [56] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *CVPR*, pages 2437–2446, 2019. 3
- [57] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. NeRV: Neural reflectance and visibility fields for relighting and view synthesis. In *CVPR*, pages 7495–7504, 2021. 3
- [58] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Generalizable patch-based neural rendering. *arXiv preprint arXiv:2207.10662*, 2022. 2
- [59] Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Light field neural rendering. In *CVPR*, pages 8269–8279, 2022. 6
- [60] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM TOG*, 38(4):1–12, 2019. 2, 3
- [61] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. In *CVPR*, pages 551–560, 2020. 3
- [62] Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. Multi-view supervision for single-view reconstruction via differentiable ray consistency. In *CVPR*, pages 2626–2634, 2017. 3
- [63] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 4, 5
- [64] 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 *CVPR*, pages 4690–4699, 2021. 2, 3, 6, 7, 8
- [65] Minye Wu, Yuehao Wang, Qiang Hu, and Jingyi Yu. Multi-view neural human rendering. In *CVPR*, pages 1682–1691, 2020. 3
- [66] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *CVPR*, pages 9421–9431, 2021. 3

- [67] Jianfeng Xiang, Jiaolong Yang, Yu Deng, and Xin Tong. Gram-hd: 3d-consistent image generation at high resolution with generative radiance manifolds. *arXiv preprint arXiv:2206.07255*, 2022. 3
- [68] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. MVSNet: Depth inference for unstructured multi-view stereo. In *ECCV*, pages 767–783, 2018. 2, 3
- [69] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *NeurIPS*, 34:4805–4815, 2021. 3
- [70] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, pages 4578–4587, 2021. 2, 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 TOG*, 40(4):1–18, 2021. 3
- [72] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. 6
- [73] Xiaoming Zhao, Fangchang Ma, David Güera, Zhile Ren, Alexander G Schwing, and Alex Colburn. Generative multiplane images: Making a 2d gan 3d-aware. *arXiv preprint arXiv:2207.10642*, 2022. 3
- [74] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: learning view synthesis using multiplane images. *ACM TOG*, 37(4):1–12, 2018. 1, 3