

# NeRFLight: Fast and Light Neural Radiance Fields using a Shared Feature Grid

Fernando Rivas-Manzanaque<sup>1,2</sup>Jorge Sierra-Acosta<sup>1</sup>Adrian Penate-Sanchez<sup>3</sup>Francesc Moreno-Noguer<sup>5</sup>Angela Ribeiro<sup>4</sup><sup>1</sup> Arquimea Research Center<sup>2</sup> Universidad Politécnica de Madrid<sup>3</sup> Universidad de las Palmas de Gran Canaria, IUSIANI<sup>4</sup> Centre for Automation and Robotics, CSIC-UPM<sup>5</sup> Institut de Robòtica i Informàtica Industrial, CSIC-UPC

## Abstract

While original Neural Radiance Fields (NeRF) have shown impressive results in modeling the appearance of a scene with compact MLP architectures, they are not able to achieve real-time rendering. This has been recently addressed by either baking the outputs of NeRF into a data structure or arranging trainable parameters in an explicit feature grid. These strategies, however, significantly increase the memory footprint of the model which prevents their deployment on bandwidth-constrained applications. In this paper, we extend the grid-based approach to achieve real-time view synthesis at more than 150 FPS using a lightweight model. Our main contribution is a novel architecture in which the density field of NeRF-based representations is split into  $N$  regions and the density is modeled using  $N$  different decoders which reuse the same feature grid. This results in a smaller grid where each feature is located in more than one spatial position, forcing them to learn a compact representation that is valid for different parts of the scene. We further reduce the size of the final model by disposing of the features symmetrically on each region, which favors feature pruning after training while also allowing smooth gradient transitions between neighboring voxels. An exhaustive evaluation demonstrates that our method achieves real-time performance and quality metrics on a par with state-of-the-art with an improvement of more than  $2\times$  in the FPS/MB ratio.

## 1. Introduction

The use of 3D objects reconstructed from real images is becoming popular in a number of applications, such as virtual-reality or online video-games. The increasing need for realistic elements makes image-based reconstruction an adequate alternative to modeling these objects from scratch using fully manual or semi-automatic design engines. At the same time, however, the larger the number of these assets is, the higher the constrain of their size in order to meet




Figure 1. NeRFLight is able to double the FPS/MB ratio of the second best method while obtaining similar quality metrics to state-of-the-art.

certain bandwidth and storage requirements.

Methods based on neural rendering have arisen as a promising approach to solve this challenge. In particular, Neural Radiance Fields (NeRF) [25] have demonstrated outstanding potential. Instead of modeling the plenoptic function [1] by means of explicit geometry representations such as point clouds [2] or voxels [22], NeRF uses a coordinate-based multi-layer perceptron (MLP) to model a density and a color field that acts as a proxy of the plenoptic function. Applying this MLP in combination with positional encoding, NeRF achieves exceptional results in scene representation, improving the quality and compactness of

the previous methods [22, 24] and resulting in an explosion of variants [4, 30, 55]. However, one of the major limitations of NeRF-based methods is their low rendering speed, being a barrier to being used in applications that require real-time.

This limitation has been addressed in different works, by either “baking” the outputs of NeRF in a data structure [10, 11, 54] or using additional features located in an explicit grid [26, 48]. Nevertheless, all these real-time NeRF rendering methods sacrifice the compactness of the representation and require relatively large models, resulting in a poor render-time performance vs storage-cost ratio (FPS/MB). This prevents these models from being deployed in applications that, besides requiring real-time, have memory storage constraints.

In this work, we introduce NeRFLight, a new approach to reduce the size of grid-based NeRF models while achieving real-time view synthesis and retaining the rendering quality of slower and computationally expensive models. Our key idea consists in splitting the volume of the density field used in NeRF into eight different regions, each with a different decoder, but sharing a common feature grid. This allows an  $8\times$  reduction in the number of total features. We also propose a novel configuration of the features that exploits the symmetry of neighboring voxel grids and favors a seamless reconstruction throughout all regions. This not only increases the rendering quality but also results in a more compact model. In addition, we also leverage the deterministic volume integration introduced in [48] to obtain a better representation than using Monte Carlo integration.

We perform an exhaustive evaluation on both synthetic and real data in terms of the FPS/MB ratio<sup>1</sup>.

The NeRFLight architecture we propose achieves rendering speeds of up to 181 FPS with models of only 14MB. As shown in Fig. 1 this yields an FPS/MB metric that is  $2\times$  larger than the closest approach in the state-of-the-art, while achieving a rendering quality (measured using PSNR and SSIM) on par with much larger models.

To summarize, our main contributions are the following:

- 1) We introduce a NeRF architecture based on several density fields and a shared feature grid that provides a light and fast representation;
- 2) We devise an approach that enforces symmetry of neighboring voxels that favors a seamless reconstruction, increases the model accuracy and further reduces its size;
- 3) All this results in the Neural Radiance Field implementation with the highest frame rate vs. storage cost ratio of the current state of the art.

## 2. Related work

**Neural Volume Reconstruction:** Novel view synthesis of 3D scenes is a well-known field in computer vision. In

<sup>1</sup>In a similar way as done in [10], where an evaluation metric of *FPS/Watts* was used to measure the compromise of *speed vs power consumption*.

the last years, research has turned towards the use of neural network representations [2, 22, 24] in what has been called Neural Rendering or Neural Volume Reconstruction. Such neural networks are applied using intermediary explicit geometry representations like point clouds [2, 46], meshes [35, 36, 45], voxel grids using CNNs [22, 33, 38] or multi-plane images [24, 41, 57].

Recently, the use of neural fields has become popular in this task. Neural fields apply coordinate-based networks to encode the relation between spatiotemporal coordinates and a given physical magnitude (see [51] for a detailed review). To achieve novel view synthesis, different types of fields (magnitudes) have been used to act as proxies of the plenoptic function [1] of a given scene. For example, [13, 39, 53] apply signed distance fields and an appearance field (color). [27, 37] replace signed distance fields by occupancy fields. Neural Radiance Fields (NeRF) [25] proposed a simple yet accurate approach that achieved unprecedented results using density and color fields. Thus, most of the subsequent works have built upon this strategy, aiming to extend original NeRF to different situations, such as unbounded scene reconstruction [3, 34, 44, 50], scene relighting [4, 5, 40, 56], scene composition [16, 21, 52], dynamic scenes [8, 17, 18, 28, 30, 47, 49] or fast convergence [9, 26, 42].

**Baking NeRF:** Several approaches have addressed the problem of speeding up NeRF at inference. The first strategy consists in “baking” the output of a pre-trained NeRF. For instance, KiloNeRF [32] uses a pretrained NeRF to train a uniform spatial grid of small MLPs in a teacher-student manner. The original NeRF is then baked into a grid of tiny NeRFs that accelerate rendering, although the use of a large number of MLPs is very memory-demanding. [54] learns first a NeRF that predicts the factors of a spherical harmonic that are embedded into a PlenOctree structure capable of modeling the view-dependent effects, so no MLP evaluation is needed at inference. However, explicitly storing spherical harmonic coefficients for the whole scene is not memory efficient. Efficient-NeRF [11] achieves the fastest rendering speed by storing the density output of NeRF and generating an octree-based structure (NeRFTree). Despite the acceleration they provide, the large memory footprint of this structure and the density grid results in a low rendering frame rate vs storage ratio. SNeRG [10] directly tackles the reduction of memory footprint. For this purpose, this method bakes the alpha composition and color along the ray together with view-based features that are decoded by a small MLP at inference. The baked information is stored in a texture atlas, where image compression algorithms such as PNG or H264 are applied. However, this compression leads to relatively low-quality metrics.

**Explicit feature grids:** The last strategy, and possibly the most popular, consists in arranging a set of trainable pa-

rameters as an explicit feature grid, reducing the number of required operations at each model inference. These features are linearly interpolated to serve as an input to small and efficient MLPs. [48] leverages this interpolation to introduce a deterministic volume integration that leads to an accurate and fast-rendering model. Nevertheless, feature storage lowers the resulting FPS/MB metric. To reduce feature memory footprint, Instant-NGP [26] uses each feature at several spatial locations using a hash-grid and applies a multiresolution representation to deal with feature collision. It also applies half-precision for feature storage, improving the compactness of the model. However, the multiple linear interpolations needed due to the multiresolution hash-grid lead to a reduction in rendering speed. Finally, methods such as TensoRF [6], or VQAD [43] apply more elaborated feature transformations like tensor decomposition or vector quantization, respectively. These strategies produce very compact models, however, they also make feature retrieval more complex which prevents real-time rendering.

NeRFLight also relies on explicit feature grids. As in [48] we also apply deterministic integration to accelerate rendering and improve the accuracy of the reconstruction. On the other hand, we also apply features located at different regions of the space. However, instead of using a multi-resolution scheme, we apply multiple density decoders (like [31,32]) to deal with feature collisions, avoiding the rendering speed loss of [26]. We do not make use of a hash grid. Instead, we have devised a more optimal symmetric configuration for the features that leads to high-quality seamless reconstruction and even faster and lighter models.

### 3. Background

NeRF [25] and its variants represent 3D scenes by approximating the plenoptic function as a density field  $\sigma(\mathbf{x})$  and a color field  $\mathbf{c}(\mathbf{x}, \mathbf{d})$ , which vary depending on the spatial position  $\mathbf{x}$  and viewing direction  $\mathbf{d}$ . These fields are encoded using a coordinate-based MLP with weights  $\mathbf{w}$ . Radiance is accumulated along a ray  $\mathbf{r}(t) = \mathbf{o} + \mathbf{d}t$  using the volume rendering equation [12]:

$$\hat{\mathbf{c}}(\mathbf{r}) = \int_0^\infty e^{-\int_0^t \sigma(\mathbf{r}(\tau))d\tau} \sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt, \quad (1)$$

to obtain the final color of a given pixel. Eq. (1) is approximated using Monte Carlo integration on a set of colors and densities coming from the inference of the MLP at  $n$  randomly sampled points along the ray ( $\mathbf{c}_i, \sigma_i = \text{MLP}_{\mathbf{w}}(\mathbf{x}_i, \mathbf{d}_i)$ ). Radiance and density are assumed to be constant between samples of a single ray. Thus, Eq. (1) can be approximated as:

$$\hat{\mathbf{c}}(\mathbf{r}) = \sum_{i=1}^n T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i \quad (2)$$

$$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right), \quad (3)$$

where  $T_i$  corresponds to the accumulated transmittance at a given point and  $\delta_i$  to the length of the interval between points  $\delta_i = \|\mathbf{x}_{i+1} - \mathbf{x}_i\|_2$ . NeRF-like models learn to approximate the density and color fields of the scene by minimizing the photometric loss between the pixels rendered using Eq. (2) and the pixels from the training images.

In order to accelerate rendering, some methods [20,48] arrange additional training parameters as explicit features located at the vertices of a voxel grid that are used as input for the MLP. This allows to reduce its size and, since the number of parameters involved in each model inference is also reduced, a faster rendering speed is obtained. To achieve this, these features are linearly interpolated depending on the position  $\mathbf{x}$  of the sampled point:

$$(\sigma_i, \mathbf{c}_i) = \text{MLP}_{\mathbf{w}}\left(\hat{\mathbf{f}}(\mathbf{x}), \mathbf{d}\right), \quad (4)$$

where  $\hat{\mathbf{f}}(\cdot)$  is the trilinear interpolation of the features of a given voxel. Despite of the achieved acceleration, when using explicit feature grids, most of the compactness provided by NeRF's MLP is lost due to the storage required cost for the feature grid. This has a negative impact for the FPS/MB metric.

## 4. Method

As mentioned in Sec. 3, in order to reduce computational cost we represent a 3D scene as a density and a color field. Similar to [48], we apply deterministic volume integration and implicit feature initialization. We also use a coarse-to-fine strategy to speed up training and perform empty space pruning to accelerate model inference. However, unlike other grid-based methods, our representation is compact thanks to the proposed shared feature grid we describe below.

### 4.1. Shared Feature Grid

As shown in Fig. 2, the scene is split into  $N$  different regions. NeRFLight makes use of a shared feature grid that is repeated in each of them. Therefore, the features are arranged at several locations, forcing the information to be more compact as it needs to represent more than one part of the scene. At a given location, a region-specific density decoder  $\theta_i$  is used to convert the corresponding features into the density value  $\sigma$  and an intermediate representation  $h$ . Finally, a single-color decoder uses this intermediate representation to obtain the color as a function of the viewing direction. Therefore, for a given point  $\mathbf{x}_j$  in a ray  $\mathbf{r}_j(t) = \mathbf{o}_j + \mathbf{d}_j t$  we have:

$$\sigma_j, h_j = \theta_i(\hat{\mathbf{f}}(\mathbf{x}_j)), \quad i \in [1, N], \quad (5a)$$




Figure 2. **NeRFLight overview.** Our method represents a scene as a set of  $N$  density fields and a single color field parameterized by  $N$  density decoders  $\theta_i$ , a color decoder  $\theta_{color}$  alongside an explicit feature grid. This feature grid is repeated for each of the regions, i.e., density fields, and is shared by all the decoders. The features are placed symmetrically to the scene center on each volume, which favours continuity between density fields and feature pruning after training.

$$\mathbf{c}_j = \theta_{color}(h_j, \gamma(\mathbf{d}_j)), \quad (5b)$$

where  $\hat{\mathbf{f}}(\cdot)$  is the integrated feature we describe later in Sec. 4.3 and  $\gamma(\mathbf{d}_j)$  is the positional encoding of the viewing direction, as in [25]. Note that Eq. (5a) is equivalent to representing the scene using eight different density fields while, as depicted by Eq. (5b), a single color field is used. This representation allows for reducing the number of unique features in the voxel grid and, subsequently, the memory needed by the feature grid, by a factor of  $N$ . As a design choice, we consider each of the  $N$  regions to be a cube with a size equal to half the side of the bounding box, resulting in a total of  $N = 8$  regions and their corresponding density decoders. For all the regions, the voxel grid resolution is kept constant and it is chosen independently for each dataset. The dimension of the features located at each vertex of these grids is set to 32.

## 4.2. Symmetric Voxel Grid

A natural approach when arranging the features in the voxel grid would be to simply repeat them in the same order for each region of the scene, using a linear configuration as shown Fig. 3a. However, this may introduce discontinuities at the seams between regions. Some previous works using neural fields solve this discontinuity issue by overlapping the different regions [23]. While this might be effective in a fully-implicit representation, doing this in an explicit grid of  $n^3$  features would increase the number of features by  $3n^2 + 3n + 1$ , resulting in a growth of the model size and reducing the effectiveness of our approach. Instead, we propose not

to modify the regions of the scene but to arrange the features symmetrically.

Let us consider a voxel vertex  $\mathbf{v}$  in a linear voxel grid (Fig. 3a), located at the boundary of regions  $e_1$  and  $e_2$ . Note that in this case, the feature that will be seen by decoder  $\theta_1$  and decoder  $\theta_2$  at  $\mathbf{v}$  will not be the same, discouraging the model from producing coherent results at both sides of the seam. In contrast, if the features are arranged symmetrically w.r.t. the scene center (Fig. 3b), we force the feature seen by any decoder at any boundary to be the same as the feature seen by all its neighboring decoders.

In practice, we also initialize all the density decoders with the same weights to improve continuity during training. The symmetric configuration also increases the probability of a feature being on empty space for all the regions of the volume. Therefore, this approach not only avoids increasing the number of features needed for a seamless reconstruction, but it also reduces the number of features that have to be stored after the pruning stage, further reducing the model size (see again Fig. 3). In Sec. 5.3, we will show how the use of a symmetric voxel grid turns out to be very effective to improve the model quality and reduce its size.

## 4.3. Deterministic Volume Integration

We apply deterministic integration to improve the quality of our model. Intersecting a ray with the voxel grid results in a series of intervals corresponding to the different voxels crossed by the ray. Instead of randomly sampling a point and then applying trilinear interpolation of the features of




Figure 3. **Symmetric Voxel Grid configuration.** Arranging the features in a symmetric configuration within the voxel grid favours the continuity between regions by forcing the different decoders to share the same features at the boundaries. It also reduces the number of features needed to represent the scene, reducing the model size.

the given voxel to obtain the input feature, the trilinear interpolation operation is integrated along each of the intervals:

$$(\sigma_j, h_j) = \theta_i \left( \int_{t_j^{\text{in}}}^{t_j^{\text{out}}} \hat{\mathbf{f}}(\mathbf{x}_j) dt, \mathbf{d} \right), \quad i \in [1, 8], \quad (6a)$$

$$\mathbf{c}_j = \theta_{color}(h_j, \gamma(\mathbf{d}_j)), \quad (6b)$$

where  $t_j^{\text{in}}$  is the ray entry point of each voxel and  $t_j^{\text{out}}$  is the ray exit point of each voxel. This deterministic feature integration produces a better estimation for Eq. (1) compared with Monte Carlo since the density and color fields are not considered to be constant between samples. Instead, the whole interval is represented by the integration of the features belonging to the voxel where the given interval is defined.

#### 4.4. Architecture and model training

The density decoders are composed of an input layer and a hidden layer of 32 neurons and an output layer of 33 neurons, while the color decoder consists only of an input layer of 32 neurons and an output layer of three neurons. All layers are activated with ReLU except for output layers, which have no activation in the case of density decoders and use sigmoid for the color decoder. For the implicit MLP, we use an architecture consisting of one input layer and eight hidden layers with 512 neurons, and an output layer with 32 neurons. All layers use ReLU activation except for the last one, where no activation is used. The positional encoding of the vertex coordinates is concatenated to the output of the fourth hidden layer.

Before training our model, we perform coarse geometry searching using a low-resolution voxel grid with a single region and one single-density decoder. We train this coarse model for five epochs and extract an occupancy mask, discarding the weights and the features of the model. This occupancy mask is obtained in a completely unsupervised

way, i.e. it is extracted from the density outputs of the model. We use this mask to speed up our training by avoiding optimizing features located in empty space.

As discussed in [48], deterministic integration is prone to overfit at certain voxels when features are explicitly trained (i.e. directly optimized). To overcome this issue, we adopt the same hybrid regularization of the features proposed by [48]. This hybrid approach is composed of two stages: in the first one, features are implicitly optimized (i.e. an MLP is trained to obtain the features given the vertex coordinates). In the second one, the implicit MLP is discarded and features are explicitly optimized, leading to a better final quality. In this way, the MLP used at the implicit stage provides a powerful initialization that avoids overfitting.

To optimize our model, besides the photometric loss, we apply the sparsity regularization loss of [10]:

$$L_{\text{sparsity}} = \lambda_s \sum_i \log \left( 1 + \frac{\sigma_i^2}{0.5} \right), \quad (7)$$

where  $\sigma_i$  represents the  $i$ th accumulated density and  $\lambda_s$  is the regularization weight. This regularization favors the sparsity of the voxel grid and prevents the model from predicting background color in empty space.

#### 4.5. Model sparsification and real-time rendering

We leverage the sparsity of the voxel grid to store only the indices and values for the features of the occupied voxels. We also store a binary occupancy mask and the weights for the density and color decoders. To further optimize the model size, we follow the same strategy as Instant-NGP [26], and store the parameters using half precision.

At inference, we make use of an octree generated from the occupancy mask to avoid querying empty or occluded voxels, as in [54]. We do that by rendering the alpha maps for all training views and recording the maximum blended alpha for each voxel. Then, we prune the voxels with an

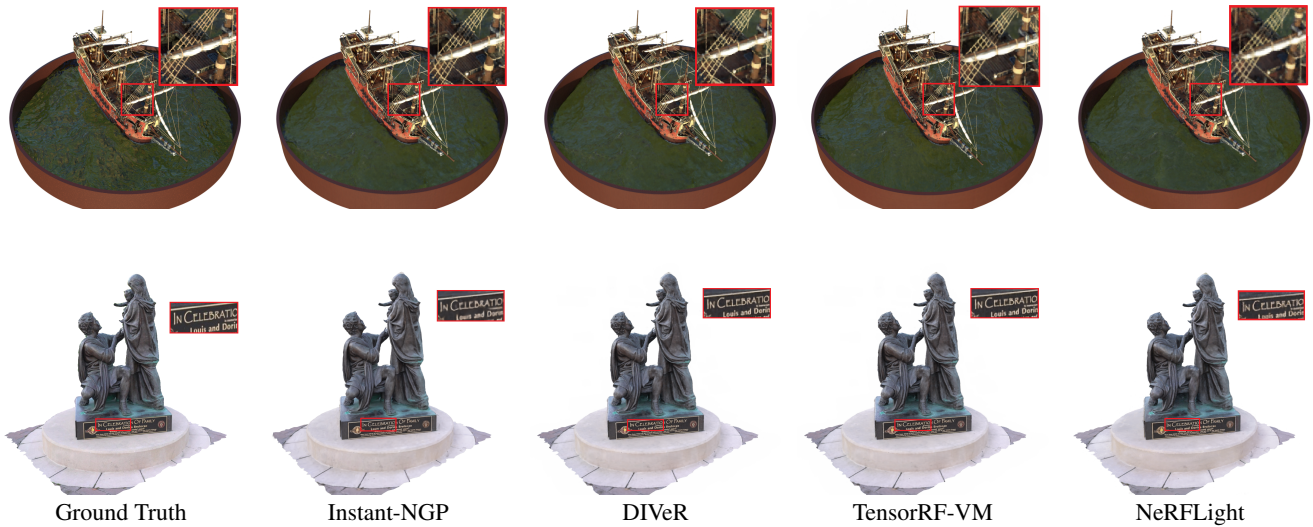


Figure 4. **Qualitative results** show that NeRFLight is able to model fine details at a similar quality than current state of the art models while achieving a  $6\times$  increment in the frame rate vs storage ratio of Instant-NGP and DIVER.

alpha value below a threshold  $\tau_\alpha$ . While rendering a ray, we stop when the accumulated transmittance is below a certain threshold  $\tau_t$ . If the alpha value of a given interval, after decoding the density value, is below  $\tau_t$ , we do not evaluate the color decoder in that interval.

## 5. Experiments

We evaluate NeRFLight and compare it with the current state-of-the-art methods on the two most established benchmarks for bounded  $360^\circ$  scenes: the NeRF-synthetic dataset [25], consisting of  $800 \times 800$  images of eight synthetic scenes and a subset of the Tanks and Temples dataset [15], provided in [20], composed of five real scenes of  $1920 \times 1080$  images where the background has been masked. The NeRF-synthetic dataset allows testing the method on scenes with complex geometry and non-Lambertian materials. On the other side, using the Tanks and Temples dataset we can evaluate our model in real-world scenes. We also analyze the contribution of the training strategy and the symmetric grid configuration on an ablation study performed on the NeRF-synthetic dataset.

### 5.1. Implementation details

We implement our method using Pytorch [29] alongside custom CUDA kernels for the ray-voxel intersection. At inference, decoders evaluation and alpha blending are also implemented using CUDA kernels. The voxel grid resolution is set to  $256^3$  for the NeRF-synthetic dataset (256 voxels per side of the cubic voxel grid) and  $320^3$  for the Tanks and Temples dataset. As detailed in Sec. 4.1, the number

of voxels for each of the eight regions are  $128^3$  and  $160^3$  respectively. The resolution for the voxel grid at the coarse stage is  $1/4$  of the resolution at the fine stage, and we also downsample training images to  $1/4$  of the original resolution. We set the sparsity regularization weight to  $\lambda_s = 1e^{-5}$  and the alpha threshold for voxel pruning to  $\tau_\alpha = 0.01$ . We prune twice, after the coarse stage and after the fine stage. For rendering, we set  $\tau_t$  also to 0.01.

At training, we sample rays from the training set using a batch size of 16384 for the NeRF-synthetic dataset and 8192 for Tanks and Temples. After the coarse geometry searching, we train NeRFLight with the implicit MLP for 1000 epochs, and, then, we train using the explicit grid for other 1000 epochs. We use Adam [14] as optimizer with a learning rate of  $1e-3$  for the coarse geometry searching and  $5e-4$  for the fine stage. For our experiments, we use an NVIDIA A100 GPU during training, and the peak memory usage is 40 GB approximately. At inference, we use an NVIDIA RTX 3090 GPU, except for the method marked as \* in Tab. 1, whose metrics are the ones reported in [10] using an NVIDIA V100 GPU. We did not re-run these methods due to their extremely low FPS/MB metric.

### 5.2. Baseline Comparisons

Tab. 1 shows an exhaustive quantitative comparison of NeRFLight and current state-of-the-art methods. The quality metrics shown for the methods marked with \* for the NeRF Synthetic dataset are the ones reported in [10] while, for the Tanks and Temples dataset, are the ones reported in [32]. The rest of the quality metrics are the ones reported in the original works.

NeRF Synthetic Dataset [25]						
Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MB $\downarrow$	FPS $\uparrow$	FPS/MB $\uparrow$
NeRF* [25]	31.00	0.947	0.081	5	0.03	0.006
JAXNeRF+* [7]	33.00	0.962	0.038	18	0.01	0.001
JAXNeRF* [7]	31.65	0.952	0.051	<u>4.8</u>	0.05	0.010
AutoInt* [19]	25.55	0.911	0.170	5	0.38	0.076
SNeRG(PNG) [10]	30.38	0.950	0.050	84	<u>202</u>	2.404
SNeRG(H264) [10]	29.86	0.938	0.065	30.2	<u>202</u>	<u>6.689</u>
Eff-NeRF [11]	31.68	0.954	<b>0.020</b>	648	<b>238</b>	0.367
KiloNeRF [32]	31.00	0.950	<u>0.030</u>	161	79	0.490
Plenotrees [54]	31.71	0.958	0.053	1930	168	0.087
TensoRF-CP [6]	31.56	0.949	0.076	<b>3.9</b>	0.45	0.115
TensoRF-VM-192-30k [6]	<u>33.14</u>	0.963	0.047	71.8	0.38	0.005
Instant-NGP [26]	<b>33.18</b>	<b>0.974</b>	0.043	27	62	2.296
DIVeR [48]	32.12	0.958	0.033	68	133	1.956
NeRFLight	31.41	<u>0.968</u>	0.039	14	181	<b>12.929</b>

Tanks and Temples Dataset [15]						
Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MB $\downarrow$	FPS $\uparrow$	FPS/MB $\uparrow$
NeRF* [25]	28.32	0.900	0.110	<u>5</u>	-	-
KiloNeRF [32]	<u>28.41</u>	0.910	0.090	-	41.3	-
Plenotrees [54]	27.99	0.917	0.131	2629	42.22	0.016
TensoRF-CP [6]	27.59	0.897	0.181	<b>4.23</b>	0.27	0.064
TensoRF-VM-192-30k [6]	<b>28.56</b>	0.920	0.140	71.4	0.16	0.002
Instant-NGP [26]	27.50	<u>0.935</u>	<b>0.076</b>	40	29	<u>0.725</u>
DIVeR [48]	28.21	0.906	<u>0.082</u>	116	<u>54</u>	0.470
NeRFLight	27.85	<b>0.939</b>	0.084	40	<b>78</b>	<b>1.950</b>

Table 1. **Quantitative results** on NeRF Synthetic and Tanks and Temples shows how NeRFLight is the method with the higher frame rate vs storage ratio while achieving competitive quality metrics. In **bold** we highlight the best performing model and the underlined results indicate the second best. \* indicates the metrics are reported in a NVIDIA V100 GPU.

**NeRF Synthetic.** The results in the NeRF synthetic dataset show that current NeRF methods are not able to achieve a frame rate vs storage ratio higher than 2.5 FPS/MB with high-quality metrics. Whereas SNeRG(H264) [10] is able to achieve 6.689 FPS/MB thanks to the texture atlas compression, its quality metrics are substantially lower than models like Instant-NGP [26] or TensoRF [6]. In contrast, NeRFLight achieves 12.929 FPS/MB, almost  $2\times$  the frame rate vs storage ratio of SNeRG(H264) and more than  $5\times$  the ratio of Instant-NGP or DIVeR [48] with comparable quality metrics, achieving the second best result in SSIM. Fig. 4 shows qualitatively that NeRFLight is able to render complex fine structures at the same level of detail as the best state-of-the-art models.

**Tanks and Temples.** Comparisons on the Tanks and Temples dataset show that NeRFLight is not only the method with the highest FPS/MB metric but also the most efficient at inference. Furthermore, the quality metrics demonstrate that NeRFLight is able to achieve accurate reconstruction, obtaining the best SSIM result. Fig. 4 shows that NeRFLight is also able to model fine high-frequency details even when applied to real-world scenes.

Instant-NGP [26] reports no results on this dataset, so we used the implementation provided by the original work to train the models for this experiment. Since Instant-NGP

is highly configurable, we tune its hyperparameters so that the model size matches the size of NeRFLight. In this way, we can compare NeRFLight with the most accurate model in the current state of the art at the same model size. It can be seen how NeRFLight is able to outperform Instant-NGP in PSNR and SSIM while being able to render more than  $2\times$  faster. See supplemental materials for more details about Instant-NGP hyperparameters.

### 5.3. Ablation Study

We next conduct an ablation study of NeRFLight through two experiments that analyze the contribution of the symmetric grid and the implicit initialization of the features. First, we analyze the quality and continuity of the intermediate representations obtained by the density decoders using the Lego Bulldozer scene as an example. Then, we show quantitative and qualitative results of the different configurations of NeRFLight.

**Decoders’ continuity.** In this experiment, we analyze the contribution of the symmetric voxel grid by inspecting the continuity of the intermediate representations obtained by the density decoders ( $h_j$  in Eq. (5b)). To do this, we apply alpha-composition of the intermediate representation at each point and apply PCA to reduce their dimension from 32 to 3, so they can be rendered as an image. Fig. 5 shows the results on the Lego scene. It can be seen how applying a




Figure 5. **PCA decomposition** of density decoder’s outputs shows that a symmetric grid is able to remove seams while a linear grid produces discontinuities.

Grid	Reg.	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FPS/MB $\uparrow$
Linear	Im	29.82	0.947	0.050	172/24=7.17
	Im-Ex	<u>30.98</u>	0.956	<u>0.046</u>	172/24=7.17
Symm	Im	30.56	<u>0.959</u>	<u>0.046</u>	<b>181/14=12.9</b>
	Im-Ex	<b>31.41</b>	<b>0.968</b>	<b>0.039</b>	<b>181/14=12.9</b>

Table 2. **Ablation of NeRFLight** shows that symmetric voxel grid configuration and implicit feature initialization not only improve the quality of the representation, but also improve the frame rate vs storage ratio.

symmetric grid produces not only a seamless representation but also a better-clustered feature space. On the other side, applying a linear grid produces discontinuities, so the color decoder will have to struggle with this inconsistent representation. To demonstrate that our strategy is also robust to non-symmetric scenes, we provide another version of Fig. 5 in a rotated Lego scene in the supplemental materials.

**Quality of the reconstruction.** The artifacts produced at the frontiers between neighboring regions when applying a linear non-symmetric grid can be clearly seen in Fig. 6a. Fig. 6b shows that a symmetric grid is able to overcome this problem by producing smooth seamless transitions. Tab. 2 reports the quality and FPS/MB metrics for the different configurations of NeRFLight. These results demonstrate that applying a symmetric grid not only improves the quality of the reconstruction but also helps to achieve a higher frame rate vs storage ratio by reducing the number of features. This leads to a lighter model and faster access to the features for each voxel (faster rendering). Finally, Tab. 2 also shows that by combining implicit and explicit training, the quality of the final model is enhanced. We also provide some results applying direct optimization of the features (i.e. fully explicit regularization) in the supplemental materials, where we show that our method also achieves high-quality results using this training method.

## 6. Limitations and future work

The most important limitation of NeRFLight is the amount of training time it requires, which is much larger




Figure 6. **Continuity at seams.** Locating the features symmetrically to the scene center in the voxel grid removes discontinuities at the seams between regions.

than other methods like TensorRF [6] or Instant-NGP [26]. However, we believe that applying tensor decomposition as in TensorRF would reduce the convergence time of NeRFLight. Another limitation of our current implementation is that it is not able to represent unbounded or forward-facing scenes. In future work, we aim to address these other types of scenes by applying normalized device coordinates (NDC) or multi-sphere images (MSI). Finally, we believe that related research approaches might benefit from applying shared feature grids to their grid-based representations.

## 7. Conclusions

In this work, we have presented NeRFLight, the method with the highest frame rate vs storage cost ratio of the state-of-the-art. We show how, even when optimizing for reduced size and inference efficiency, our method provides high-quality rendering results, on par with the most recent NeRF models. We also introduced the concept of a shared feature grid in NeRF, as a way to reuse features to make grid representations more compact. Furthermore, we have proposed the first approach to achieve seamless reconstructions using symmetric voxel grids. We have shown that this strategy, reduces the number of features in the grid and leads to a faster rendering and a more accurate reconstruction. To summarize, NeRFLight is the current NeRF-based method with the best FPS/MB ratio, which makes it especially adequate to stream and render in scenarios like social media, online games, or for upcoming metaverse applications.

## Acknowledgements

This work is partially supported by the European Union under the project FlexiGroBots (DOI: 10.3030/101017111) and by the Spanish government under projects MoHuCo (PID2020-120049RB-I00) and SmartWEED-DARWEEM (PID2020-113229RB-C43). Adrian Penate-Sanchez is supported by a Beatriz Galindo grant. The authors would like to thank Marcos Barrios for his support in the graphic design of the figures and videos included in the paper.



## References

- [1] Edward H. Adelson and James R. Bergen. The plenoptic function and the elements of early vision. In *Computational Models of Visual Processing*, pages 3–20. MIT Press, 1991. [1](#), [2](#)
- [2] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *Computer Vision – ECCV 2020*, pages 696–712, 2020. [1](#), [2](#)
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5470–5479, June 2022. [2](#)
- [4] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P.A. Lensch. Nerf: Neural reflectance decomposition from image collections. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12664–12674, 2021. [2](#)
- [5] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik PA Lensch. Neural-pil: Neural pre-integrated lighting for reflectance decomposition. In *Advances in Neural Information Processing Systems*, volume 34, pages 10691–10704, 2021. [2](#)
- [6] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision (ECCV)*, 2022. [3](#), [7](#), [8](#)
- [7] Boyang Deng, Jonathan T. Barron, and Pratul P. Srinivasan. JaxNeRF: an efficient JAX implementation of NeRF, 2020. [7](#)
- [8] Yilun Du, Yanan Zhang, Hong-Xing Yu, Joshua B. Tenenbaum, and Jiajun Wu. Neural radiance flow for 4d view synthesis and video processing. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14304–14314, 2021. [2](#)
- [9] 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 (CVPR)*, pages 5501–5510, June 2022. [2](#)
- [10] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5855–5864, 2021. [2](#), [5](#), [6](#), [7](#)
- [11] Tao Hu, Shu Liu, Yilun Chen, Tiancheng Shen, and Jiaya Jia. Efficientnerf efficient neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12902–12911, June 2022. [2](#), [7](#)
- [12] James T. Kajiya and Brian P Von Herzen. Ray tracing volume densities. *SIGGRAPH Comput. Graph.*, 18(3):165–174, jan 1984. [3](#)
- [13] Petr Kellnhofer, Lars C. Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4285–4295, 2021. [2](#)
- [14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [6](#)
- [15] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Trans. Graph.*, 36(4), jul 2017. [6](#), [7](#)
- [16] Abhijit Kundu, Kyle Genova, Xiaoqi Yin, Alireza Fathi, Caroline Pantofaru, Leonidas Guibas, Andrea Tagliasacchi, Frank Dellaert, and Thomas Funkhouser. Panoptic Neural Fields: A Semantic Object-Aware Neural Scene Representation. In *CVPR*, 2022. [2](#)
- [17] Tianye Li, Mira Slavcheva, Michael Zollhöfer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, and Zhaoyang Lv. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5521–5531, June 2022. [2](#)
- [18] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6498–6508, June 2021. [2](#)
- [19] David B. Lindell, Julien N. P. Martel, and Gordon Wetzstein. AutoInt: Automatic integration for fast neural volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14556–14565, June 2021. [7](#)
- [20] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 2020. [3](#), [6](#)
- [21] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. [2](#)
- [22] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 38(4), jul 2019. [1](#), [2](#)
- [23] Ishit Mehta, Michaël Gharbi, Connelly Barnes, Eli Shechtman, Ravi Ramamoorthi, and Manmohan Chandraker. Modulated periodic activations for generalizable local functional representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14214–14223, October 2021. [4](#)
- [24] 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 Trans. Graph.*, 38(4), jul 2019. [2](#)
- [25] 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*, 2020. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [26] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multires-

- olution hash encoding. *ACM Trans. Graph.*, 41(4), jul 2022. [2](#), [3](#), [5](#), [7](#), [8](#)
- [27] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3501–3512, 2020. [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 *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5845–5854, 2021. [2](#)
- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Proceedings of Advances in Neural Information Processing Systems*, volume 32, 2019. [6](#)
- [30] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10313–10322, 2021. [2](#)
- [31] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. Derf: Decomposed radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14153–14161, June 2021. [3](#)
- [32] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14315–14325, 2021. [2](#), [3](#), [6](#), [7](#)
- [33] Konstantinos Rematas and Vittorio Ferrari. Neural voxel renderer: Learning an accurate and controllable rendering tool. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5416–5426, 2020. [2](#)
- [34] Konstantinos Rematas, Andrew Liu, Pratul P. Srinivasan, Jonathan T. Barron, Andrea Tagliasacchi, Tom Funkhouser, and Vittorio Ferrari. Urban radiance fields. *CVPR*, 2022. [2](#)
- [35] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *Computer Vision – ECCV 2020*, pages 623–640, 2020. [2](#)
- [36] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12216–12225, June 2021. [2](#)
- [37] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Hao Li, and Angjoo Kanazawa. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2304–2314, 2019. [2](#)
- [38] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Niessner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#)
- [39] Vincent Sitzmann, Michael Zollhofer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. [2](#)
- [40] 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 *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7491–7500, 2021. [2](#)
- [41] Pratul P. Srinivasan, Richard Tucker, Jonathan T. Barron, Ravi Ramamoorthi, Ren Ng, and Noah Snavely. Pushing the boundaries of view extrapolation with multiplane images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [2](#)
- [42] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5459–5469, June 2022. [2](#)
- [43] Towaki Takikawa, Alex Evans, Jonathan Tremblay, Thomas Müller, Morgan McGuire, Alec Jacobson, and Sanja Fidler. Variable bitrate neural fields. In *ACM SIGGRAPH 2022 Conference Proceedings*. Association for Computing Machinery, 2022. [3](#)
- [44] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul Srinivasan, Jonathan T. Barron, and Henrik Kretzschmar. Block-NeRF: Scalable large scene neural view synthesis. *arXiv*, 2022. [2](#)
- [45] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Trans. Graph.*, 38(4), jul 2019. [2](#)
- [46] Cen Wang, Minye Wu, Ziyu Wang, Liao Wang, Hao Sheng, and Jingyi Yu. Neural opacity point cloud. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(7):1570–1581, 2020. [2](#)
- [47] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Fourier plenotrees for dynamic radiance field rendering in real-time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13524–13534, June 2022. [2](#)
- [48] Liwen Wu, Jae Yong Lee, Anand Bhattad, Yu-Xiong Wang, and David Forsyth. Diver: Real-time and accurate neural radiance fields with deterministic integration for volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16200–16209, June 2022. [2](#), [3](#), [5](#), [7](#)
- [49] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9421–9431, 2021. [2](#)

- [50] Yuanbo Xiangli, Linning Xu, Xingang Pan, Nanxuan Zhao, Anyi Rao, Christian Theobalt, Bo Dai, and Dahua Lin. Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. In *The European Conference on Computer Vision (ECCV)*, 2022. [2](#)
- [51] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *Computer Graphics Forum*, 2022. [2](#)
- [52] 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 *International Conference on Computer Vision (ICCV)*, October 2021. [2](#)
- [53] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 2492–2502. Curran Associates, Inc., 2020. [2](#)
- [54] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5732–5741, 2021. [2](#), [5](#), [7](#)
- [55] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020. [2](#)
- [56] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul Debevec, William T. Freeman, and Jonathan T. Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graph.*, 40(6), dec 2021. [2](#)
- [57] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM Trans. Graph.*, 37(4), jul 2018. [2](#)