

# Learning Neural Light Fields with Ray-Space Embedding

Benjamin Attal\*  
Carnegie Mellon University

Jia-Bin Huang  
Meta

Michael Zollhöfer  
Reality Labs Research

Johannes Kopf  
Meta

Changil Kim  
Meta

<https://neural-light-fields.github.io>

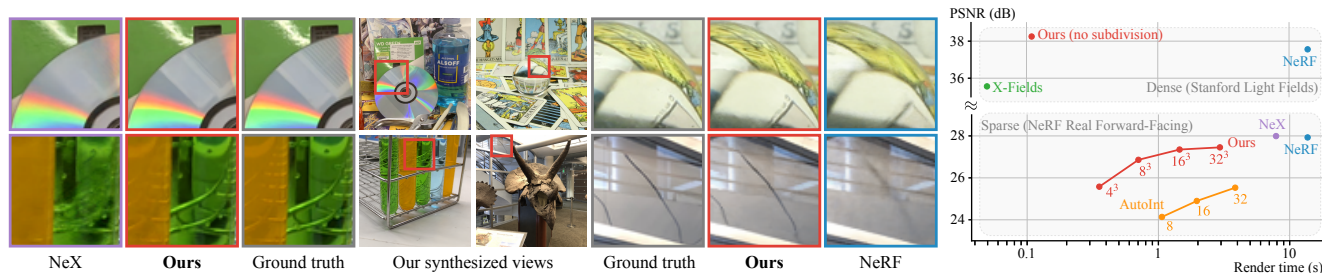


Figure 1. We present *Neural Light Fields with Ray-Space Embedding*, which map rays directly to integrated radiance. Our representation allows for high-quality view synthesis with faithful reconstruction of complex view dependence (*left*). We are able to synthesize novel views in a fraction of the time required for the current state-of-the-art approaches, with comparable memory footprint. Our method achieves a more favorable trade-off between quality, speed, and memory than existing approaches for both dense and sparse light fields (*right*). See more visual comparisons including video results in our *project website* linked above.

## Abstract

*Neural radiance fields (NeRFs) produce state-of-the-art view synthesis results, but are slow to render, requiring hundreds of network evaluations per pixel to approximate a volume rendering integral. Baking NeRFs into explicit data structures enables efficient rendering, but results in large memory footprints and, in some cases, quality reduction. Additionally, volumetric representations for view synthesis often struggle to represent challenging view dependent effects such as distorted reflections and refractions. We present a novel neural light field representation that, in contrast to prior work, is fast, memory efficient, and excels at modeling complicated view dependence. Our method supports rendering with a single network evaluation per pixel for small baseline light fields and with only a few evaluations per pixel for light fields with larger baselines. At the core of our approach is a ray-space embedding network that maps 4D ray-space into an intermediate, interpolable latent space. Our method achieves state-of-the-art quality on dense forward-facing datasets such as the Stanford Light Field dataset. In addition, for forward-facing scenes with sparser inputs we achieve results that are competitive with NeRF-based approaches while providing a better speed/quality/memory trade-off with far fewer network evaluations.*

\* This work was done while Benjamin was an intern at Meta.

## 1. Introduction

View synthesis is an important problem in computer vision and graphics. Its goal is to photorealistically render a scene from unobserved camera poses, given a few posed input images. Existing approaches solve this problem by optimizing some underlying representation of the scene’s appearance and geometry and then rendering this representation from novel views.

View synthesis has recently experienced a renaissance with an explosion of interest in *neural scene representations* — see Tewari et al. [52,53] for a snapshot of the field. *Neural radiance fields* [30] are perhaps the most popular of these neural representations, and methods utilizing them have recently set the state-of-the-art in rendering quality for view synthesis. A radiance field is a 5D function that maps a 3D point  $\mathbf{x}$  and 3D direction  $\vec{\omega}$  (with only 2 degrees of freedom) to the radiance leaving  $\mathbf{x}$  in direction  $\vec{\omega}$ , as well as the density of the volume at point  $\mathbf{x}$ . A neural radiance field or NeRF represents this function with a neural network. Because volume rendering a NeRF is differentiable, it is straightforward to optimize by minimizing the difference between ground truth views at known camera poses and their reconstructions.

The main drawback of neural radiance fields is that volume rendering requires many samples and thus many neural network evaluations per ray to approximate a volume rendering accurately. Thus, rendering from a NeRF is usually quite slow. Various approaches exist for baking or caching neural

radiance fields into explicit data structures to improve efficiency [11, 15, 39, 66]. Some approaches, concurrent to this work, learn color and density directly on a voxel grid, which improves both training and rendering speed [31, 50, 65]. However, the storage cost for explicit representations is much higher than a NeRF. Further, the baking procedure itself sometimes leads to a loss in resulting view synthesis quality for baking methods. Other methods aim to reduce the number of neural network evaluations per ray by representing radiance only on surfaces [17, 32]. These methods predict new images with only a few evaluations per ray. However, their quality is contingent on either ground truth geometry or reasonable geometry estimates, which are not always available.

A *light field* [12, 20] is the integral of a radiance field. It maps ray parameters directly to the integrated radiance along that ray. Thus, only one look-up of the underlying representation is required to determine the color of a ray; hence one evaluation per pixel, unlike hundreds of evaluations required by a NeRF. A common assumption for light fields is that this integral remains the same no matter the ray origin (i.e., radiance is constant along rays), which holds when the convex hull of the scene geometry does not contain any viewpoints used for rendering [20]. Given this assumption, a light field is a function of a ray on a 4D ray space.

In this paper, we show how to learn *neural light fields*. Since coordinate-based neural representations have been successfully employed to learn radiance fields from a set of ground truth images, one might think that they could be useful for representing and learning light fields as well. However, we show that learning light fields is significantly more challenging than learning radiance fields. Using the same neural network architecture as in NeRF to parameterize a light field leads to poor interpolation quality for view synthesis.

While a radiance field is a 5D function, an essential ingredient of NeRF is that it learns the scene geometry as a density field in 3D space. Additionally, a NeRF’s learned appearance is closer to a 3D function than a 5D function since the network is late-conditioned on viewing directions [67]. This makes NeRFs easy to optimize but also means that they can struggle to represent complex view-dependent effects such as reflections and refractions [59] which violate multi-view color constraints.

On the other hand, we face the problem of learning a function defined on a 4D ray space from only partial observations — input training images only cover a few 2D slices of the full 4D space. At the same time, NeRF enjoys multiple observations of most 3D points. Further, light fields do not entail any form of scene geometry, which allows them to capture complex view dependence but poses significant challenges in interpolating unseen rays in a geometrically meaningful way. Existing methods address these issues by sacrificing

view interpolation [8], leveraging data driven priors [47], or relying on strong supervision signals [22].

In order to deal with these challenges, we employ a novel *ray-space embedding network* that re-maps the input ray-space into an embedded latent space. This facilitates both the registration of rays observing same 3D points and the interpolation of unobserved rays, which leads to better memorization and view synthesis at the same time. The embedding network alone already provides state-of-the-art view synthesis quality for densely sampled inputs (such as the Stanford light fields [58]). However, it does not interpolate well in sparser input sequences (such as those from Real Forward-Facing [30]). Thus, we represent such scenes with a set of *local light fields*, where each local light field is less prone to large depth and texture changes. Each local light field has to learn a simpler embedding at the price of several more network evaluations per ray.

We evaluate our method for learning neural light fields in sparse and dense regimes, with and without subdivisions. We compare to state-of-the-art view-synthesis methods and show that our approach achieves comparable or better view synthesis quality in both regimes, in a fraction of render time that other methods require, while still maintaining their small memory footprint (Figure 1).

In summary, our contributions are:

1. A novel neural light field representation that employs a ray-space embedding network and achieves *state-of-the-art* quality for small-baseline view synthesis without any geometric constraints.
2. A subdivided neural light field representation for large baseline light fields that leads to a good trade-off in terms of the number of network queries vs. quality, which can be optimized to achieve comparable performance to NeRF [30] and NeX [59] for sparse real-world scenes.
3. Improved capture of *view-dependent appearance* in both sparse and dense regimes (e.g., complicated reflections and refractions) that existing volume-based methods such as NeRF [30] and NeX [59] struggle to represent.

## 2. Related Work

**Light Fields.** Light fields [20] or Lumigraphs [12] represent the radiance entering all positions in 3D space, in all viewing directions. By resampling the rays from the discrete set of captured images, one can query arbitrary rays in a light field and synthesize photorealistic novel views of a static scene without necessitating geometric reconstruction. While light fields have the flexibility to represent complicated, non-Lambertian appearance, this often requires high sampling rates, and thus capturing and storing an excessive number of images. Much effort has been devoted to extending light fields to sparse, unstructured input sequences [4, 7] and synthesizing (i.e. interpolating) in-between views for

such sequences [16, 43, 57, 60, 61, 64]. Very recently, neural implicit representations have been applied to modeling light fields for better memory compactness [8] and faster rendering for view synthesis [22, 47]. Like these methods, our work leverages *neural* light fields as its core representation for view synthesis. Unlike prior work, which sacrifices view interpolation [8] or leverages meta-learning for strong data-driven priors [47], we enable view synthesis *without* explicit priors, using per-scene training only.

**Representations for View Synthesis.** Various 3D representations have been developed for rendering novel views from a sparse set of captured images. Image-based rendering techniques leverage proxy geometry to warp and blend source image content into novel views [4, 6, 37, 45]. Recent advances in image-based rendering include learning-based disparity estimation [9, 16], blending [14, 40], and image synthesis via 2D CNNs [40, 41]. Other commonly used scene representations include voxels [25, 48], 3D point clouds [1, 18, 42] or camera-centric layered 3D representations, such as multi-plane images [3, 29, 55, 59, 69] or layered depth images [13, 19, 44, 56]. Rather than using discrete representations as in the above methods, a recent line of work explores neural representations to encode the scene geometry and appearance as a continuous volumetric field [2, 23, 27, 30, 49]. We refer the readers to [52] for more comprehensive discussions of these works.

Many of these methods learn a mapping from 3D points and 2D viewing directions to color/density, which can be rendered using numerical integration – thus, rendering the color of a pixel for novel views requires hundreds of MLP evaluations. Instead, our approach directly learns the mapping from rays to color, and therefore supports efficient rendering with fewer network evaluations. In some ways, our method is similar to AutoInt [22], which predicts integrated radiance along ray *segments* and uses subdivision. The accuracy of AutoInt, however, is lower than NeRF due to the approximation of the nested volume rendering integral. By contrast, our approach leads to comparable performance to NeRF.

**Coordinate-based Representations.** Coordinate-based representations have emerged as a powerful tool for overcoming the limitations of traditional discrete representations (e.g., images, meshes, voxelized volumes). The core idea is to train an MLP to map an input coordinate to the desired target value such as pixel color [26, 33, 46, 63], signed distance [5, 34], occupancy [28], volume density [30], or semantic labels [68]. Like existing coordinate-based representation approaches, our method also learns the mapping from an input coordinate (ray) to a target scene property (color). Our core contributions lie in designing a coordinate embedding network that enables plausible view synthesis.

**Learned Coordinate Embedding.** Learned coordinate transformation or embedding has been applied to extend the capability of coordinate-based representations, such as

NeRF, for handling higher dimensional interpolation problems. For example, several works tackle dynamic scene view synthesis [10, 21, 35, 54, 62] by mapping space at each time-step into a canonical frame. Others leverage coordinate embedding for modeling articulated objects such as the human body [24, 36]. Our ray-space embedding network can be viewed as *locally deforming* the input ray coordinates such that the MLP can produce plausible view interpolation while preserving faithful reconstruction of the source views.

**Subdivision.** Recent coordinate-based representations employ 3D space partitioning/subdivision either for improving the rendering efficiency [15, 39, 66] or representation accuracy [26, 38]. Our work also leverages spatial subdivision and shows improved rendering quality, particularly on scenes with large camera baselines.

### 3. Neural Light Fields

A neural radiance field [30] represents the appearance and geometry of a scene with an MLP  $F_\theta : (\mathbf{x}_t, \vec{\omega}) \rightarrow (L_e(\mathbf{x}_t, \vec{\omega}), \sigma(\mathbf{x}_t))$  with trainable weights  $\theta$ . It takes as input a 3D position  $\mathbf{x}_t$  and a viewing direction  $\vec{\omega}$ , and produces both the density  $\sigma(\mathbf{x}_t)$  at point  $\mathbf{x}_t$  and the radiance  $L_e(\mathbf{x}_t, \vec{\omega})$  emitted at point  $\mathbf{x}_t$  in direction  $\vec{\omega}$ .

One can generate views of the scene from this MLP using volume rendering:

$$L(\mathbf{x}, \vec{\omega}) = \int_{t_n}^{t_f} \underbrace{T(\mathbf{x}, \mathbf{x}_t)}_{\text{Acc. transmittance}} \underbrace{\sigma(\mathbf{x}_t)}_{\text{Density}} \underbrace{L_e(\mathbf{x}_t, \vec{\omega})}_{\text{Emitted radiance}} dt, \quad (1)$$

where  $T(\mathbf{x}, \mathbf{x}_t) = e^{-\int_{t_n}^t \sigma(\mathbf{x}_s) ds}$  describes the accumulated transmittance for light propagating from position  $\mathbf{x}$  to  $\mathbf{x}_t$ , for near and far bounds  $t \in [t_n, t_f]$  of the scene. In practice, the integral is approximated using numerical quadrature:

$$L(\mathbf{x}, \vec{\omega}) \approx \sum_{k=1}^N \hat{T}(\mathbf{x}, \mathbf{x}_k) (1 - e^{-\sigma(\mathbf{x}_k) \Delta \mathbf{x}_k}) L_e(\mathbf{x}_k, \vec{\omega}), \quad (2)$$

where  $\hat{T}(\mathbf{x}, \mathbf{x}_k) = e^{-\sum_{j=1}^k \sigma(\mathbf{x}_j) \Delta \mathbf{x}_j}$  and  $\Delta \mathbf{x}_j = \mathbf{x}_j - \mathbf{x}_{j-1}$  is the distance between adjacent samples. An accurate approximation of (1) requires many samples and thus many neural network evaluations (on the order of hundreds) per ray.

**Light Fields.** Whereas a radiance field represents the radiance emitted at each point in space, a light field represents the total *integrated radiance* traveling along a ray. In other words, Equations (1) and (2) describe the relationship between the light field  $L$  on the left-hand side and the radiance field  $(L_e, \sigma)$  on the right-hand side. By optimizing a *neural light field* representation for  $L$ , given an input ray, one can predict this integrated radiance (and therefore render the color along this ray) with only a single neural network evaluation.

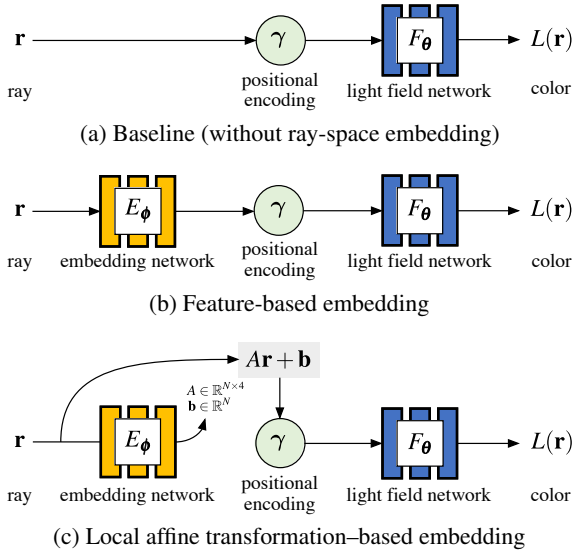


Figure 2. **Ray-space embedding networks.** Building upon the baseline (a), we explore two different schemes for ray-space embedding: (b) predicting the latent features and (c) predicting affine transformation parameters.

Assuming that we are operating on forward facing scenes, and that radiance remains constant along rays, we can construct a 4D ray space by intersecting rays with two planes,  $\pi^{xy}$  and  $\pi^{uv}$ . We can then take the local plane coordinates,  $(x, y) \in \pi^{xy}$  and  $(u, v) \in \pi^{uv}$  of these ray intersections as the ray parameterization for the 4D light field,  $L : \mathbf{r} = (x, y, u, v) \rightarrow \mathbf{c} = (r, g, b)$ .

**Baseline Neural Light Fields.** To begin with, similar to how NeRF uses an MLP to represent radiance fields, we define an MLP  $F_\theta$  to represent a light field. It takes as input positionally encoded 4D ray coordinates  $\gamma(\mathbf{r})$  and outputs color  $\mathbf{c}$  (integrated radiance) along each ray (Figure 2a):

$$L_{\text{base}}(\mathbf{r}) = F_\theta(\gamma(\mathbf{r})). \quad (3)$$

Unfortunately, this baseline approach is an unsatisfactory light field representation due to the following challenges. First, the captured input images only provide *partial observations* in the form of a sparse set of 2D slices of the full 4D ray space, so that each 4D ray coordinate in the input training data is observed at most once, assuming that no cameras share the same origin.

Second, light fields do not explicitly represent 3D scene geometry; hence the network a priori does not know how to *interpolate* the colors of unobserved rays from training observations. In other words, when querying a neural network representing a light field with *unseen* ray coordinates, multi-view consistency is not guaranteed. To address these challenges, we present two key techniques — ray-space embedding and subdivision — which substantially improve the rendering quality of the proposed neural light field representation.

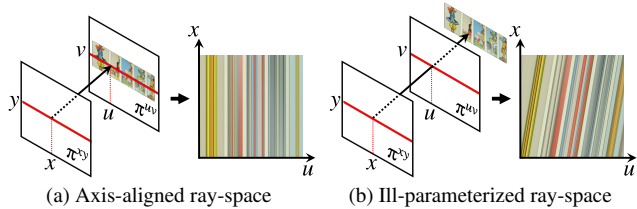


Figure 3. **Importance of light field parameterization.** (a) It is easy to learn a light field for a well-parameterized ray space, but (b) a ray space with non-axis-aligned level sets makes the problem more challenging [51]. See Section 4 and appendix for discussions.

## 4. Ray-Space Embedding Networks

We first introduce *ray-space embedding networks*. This network re-maps the input ray-space into an *embedded latent space*. Intuitively, our goals for ray-space embedding networks are:

1. **Memorization:** Map disparate coordinates in input 4D ray-space that observe the same 3D point to the same location in the latent space. This allows for more stable training and better allocation of network capacity, and thus better memorization of input views.
2. **Interpolation:** Produce an interpolable latent space, such that querying unobserved rays preserves multi-view consistency and improves view synthesis quality.

**Feature-Space Embedding.** A straightforward approach would be to learn a nonlinear mapping from 4D ray coordinates to a latent feature space via an MLP,  $E_\phi^{\text{feat}} : \mathbf{r} \rightarrow \mathbf{f} \in \mathbb{R}^N$ , where  $N$  is the feature space dimension (see Figure 2b). This embedding network can produce a nonlinear, many-to-one mapping from distinct ray coordinates into shared latent features, which would allow for better allocation of capacity in the downstream light field network  $F_\theta$ . The finite capacity of  $F_\theta$  incentivizes this “compression” and thus encourages rays with similar colors to be mapped to nearby features in the latent space; effectively a form of implicit triangulation.

After embedding, we then feed the positionally encoded  $N$ -dimensional feature into the light field MLP  $F_\theta$ :

$$L_{\text{feat}}(\mathbf{r}) = F_\theta(\gamma(E_\phi^{\text{feat}}(\mathbf{r}))). \quad (4)$$

While this approach compresses and helps to find correspondences in ray space, it is not enough to facilitate the interpolation between these correspondences (see Section 6).

**Light Field Re-parameterization.** To understand our goals better, consider a scene containing an axis-aligned textured square at a fixed  $z$ -depth. Again, we assume a two-plane parameterization for the input ray-space of this scene. If one of the planes is at the depth of the square, then the light field will only vary in two coordinates, as illustrated in Figure 3a. In other words, the light field is only a function of  $(u, v)$  and is constant in  $(x, y)$ . Since positional encoding is applied separately for each input coordinate, the baseline



model (3) can interpolate the light field perfectly if it chooses to ignore the positionally encoded coordinates  $(x, y)$ .

On the other hand, if the two planes in the parameterization do not align with the textured square, the color level sets (e.g., the line structures in the  $ux$  slices) of the light field are 2D affine subspaces that are *not* axis-aligned (see Figure 3b). Axis-aligned positional encoding, as used by NeRF [30], will yield interpolation kernels [51] that are ill-suited for capturing these sheared subspaces, especially from incomplete observations. One can therefore frame the task of effective interpolation as learning an optimal *re-parameterization* of the light field for a scene, such that the 2D color level sets for each 3D scene point are axis-aligned. See the appendix for a more extensive discussion.

**Local Affine-Transformation Embedding.** Given this intuition, we describe our architecture for ray-space embedding, which learns *local affine transformations* for each coordinate in ray space. We learn local affine transformations rather than a single global transformation because different rays may correspond to points with different depths, which will cause the shape of color level sets to vary.

To this end, we employ an MLP  $E_\phi : \mathbf{r} \rightarrow (A \in \mathbb{R}^{N \times 4}, \mathbf{b} \in \mathbb{R}^N)$ . The output of the network is a  $N \times 4$  matrix  $A$ , as well as an  $N$ -vector  $\mathbf{b}$  representing a bias (translation), which together form a  $4D \rightarrow ND$  affine transformation. This affine transformation is applied to the input ray coordinate  $\mathbf{r}$  before being positionally encoded and passed into the light field network  $F_\theta$ :

$$L(\mathbf{r}) = F_\theta(\gamma(A\mathbf{r} + \mathbf{b})), \text{ where } (A, \mathbf{b}) = E_\phi(\mathbf{r}). \quad (5)$$

See Figure 2c for an illustration of this model.

While setting  $N = 4$  above is perfectly reasonable, we use  $N = 32$  in practice, which can be seen as providing *multiple* possible re-parameterizations per ray. Also note that learning *arbitrary* affine transforms, rather than a single  $z$ -depth for each ray, allows the network to better capture both *angular frequencies* in the light field (due to object depth), as well as *spatial frequencies* (due to object texture).

## 5. Subdivided Neural Light Fields

Although this approach works well for dense light fields, our embedding network struggles to find long-range correspondences between training rays when training data is too sparse. Moreover, even if it can discover correspondences, interpolation for unobserved rays in between these correspondences remains underconstrained.

To resolve these issues, we propose learning a voxel grid of *local light fields* that covers the entire 3D scene. This approach is motivated by the following observation: if we parameterize a local light field for a voxel by intersecting rays with the voxel’s front and back planes (See Figure 4), the color level sets of the local light field are already *almost*

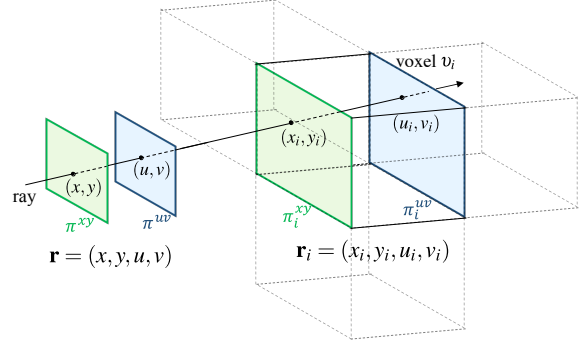


Figure 4. **Re-parameterization for local light fields.** A ray  $\mathbf{r} = (x, y, u, v)$  defined globally is re-parameterized with respect to each local light field in voxel  $v_i$ , by intersecting the front and back faces of  $v_i$ . This yields local ray coordinates  $\mathbf{r}_i = (x_i, y_i, u_i, v_i)$ .

*axis-aligned*. As such, the ray-space embedding for each voxel can be simple, only inducing small shears of the input ray-space to allow for good interpolation. On the other hand, a global light field requires learning a complex ray-space embedding that captures occlusions and significant depth changes.

Therefore a local light field can be learned much more easily than an entire global light field *with the same training data*. This still requires that we solve an assignment problem: we must know which rays to use to train each local light field. While we can exclude all rays that do not intersect the voxel containing the local light field, many other rays may be occluded before they hit the voxel. Therefore, these rays should also be excluded during training.

A simple way to handle ray assignment is by learning (per-ray) opacity. If the opacity accumulated along a ray before it hits a voxel is high, the ray should receive a little contribution from this voxel. We, therefore, modify our light field network also to produce *integrated* opacity, or alpha *for each ray*. It is important to note that the opacity can vary depending on which ray is queried through a particular voxel — opacity is a *function of the ray*, not the voxel itself.

**Subdivided Volume Rendering.** We detail how a set of local light fields can be rendered to form new images. Given a voxel grid in 3D space, we place a local light field within each voxel. A voxel’s light field is parameterized with respect to its front and back planes (See Figure 4), yielding ray coordinates  $\mathbf{r}_i$ , where  $i$  is some voxel index, for any ray  $\mathbf{r}$  defined globally. If 3D space is subdivided into  $M$  voxels, there exist  $M$  distinct parameterizations  $\{\mathbf{r}_i\}_{i=1}^M$  of each  $\mathbf{r}$ .

To implement this model, the embedding network  $E_\phi$  is augmented such that it takes the positionally encoded voxel index  $\gamma(i)$  in addition to the 4D ray coordinates, i.e.,  $E_\phi : (\mathbf{r}_i, \gamma(i)) \rightarrow (A_i, \mathbf{b}_i)$ . The light field network  $F_\theta$  is augmented in the same way and predicts both color *and* alpha:  $F_\theta : (\gamma(A_i\mathbf{r}_i + \mathbf{b}_i), \gamma(i)) \rightarrow (\mathbf{c}_i, \alpha_i)$ . Again, we emphasize that  $\alpha_i$  is a function of the embedded ray, not just the voxel.

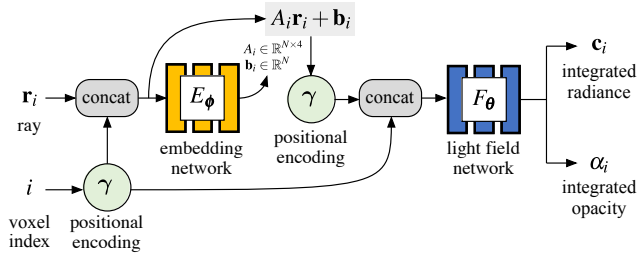


Figure 5. **Local light fields via spatial subdivision.** Each local light field takes a locally re-parameterized ray  $\mathbf{r}_i$  and its voxel index  $i$ , and predicts the radiance and opacity *both* integrated for the ray segment traveling through the voxel. Both the embedding network  $E_\phi$  and the light field network  $F_\theta$  are conditioned on positionally encoded voxel index  $\gamma(i)$ .

Given this model, rendering works in the following way. First, the set of voxels  $\mathcal{V}(\mathbf{r})$  that the ray  $\mathbf{r}$  intersects are identified. The ray is then intersected with each voxel’s front and back planes  $\pi_i^{xy}$  and  $\pi_i^{uv}$ . This process yields a series of local ray coordinates  $\{\mathbf{r}_i\}_{i \in \mathcal{V}(\mathbf{r})}$ . The color  $\mathbf{c}_i$  and alpha  $\alpha_i$  of each  $\mathbf{r}_i$  are then computed as:

$$(\mathbf{c}_i, \alpha_i) = F_\theta(\gamma(A_i \mathbf{r}_i + \mathbf{b}_i), \gamma(i)), \quad (6)$$

where the local affine-transformation is obtained from the embedding network:

$$(A_i, \mathbf{b}_i) = E_\phi(\mathbf{r}_i, \gamma(i)). \quad (7)$$

The final color of ray  $\mathbf{r}$  is then over-composited:

$$\mathbf{c} = \sum_{i \in \mathcal{V}(\mathbf{r})} \left( \prod_{j \in \mathcal{V}(\mathbf{r}) \wedge j < i} (1 - \alpha_j) \right) \alpha_i \mathbf{c}_i, \quad (8)$$

which assumes that we sort the voxels by their distance from the ray origin in ascending order (see Figure 5).

## 6. Experiments

**Datasets.** We evaluate our model in two angular sampling regimes: sparse (wide-baseline) and dense (narrow-baseline), using the Real Forward-Facing dataset [30], the Shiny dataset [59], and the Stanford Light Field dataset [58]. Each scene from the Real Forward-Facing and Shiny datasets consist of a few dozen images with relatively large angular separation in between views. We followed their authors’ train/test split and hold out every 8th image from training.

The Stanford light fields are all captured as  $17 \times 17$  dense 2D grids of images. We take every 4th image in *both* horizontal and vertical directions and use the resulting  $5 \times 5$  subsampled light field for training (thus using 25 images out of 289), while reserving the rest of the images for testing.

We downsample the images in the Real Forward-Facing and Shiny datasets to  $504 \times 378$  pixels, and use half the original resolution of the Stanford images. For quantitative metrics, we report PSNR, SSIM, LPIPS, and FPS; where FPS is computed for  $512 \times 512$  pixel images.

Table 1. **Quantitative comparisons.** We show the aggregated metrics over all scenes in each of the following datasets: Real Forward-Facing [30], *Undistorted* Real Forward-Facing [59], Shiny [59], and Stanford [58].

Dataset	Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FPS $\uparrow$
RFF	NeRF [30]	<b>27.928</b>	<b>0.916</b>	0.065	0.07
	AutoInt [22]	25.531	0.853	0.156	<u>0.26</u>
	Ours (subdivision)	<u>27.454</u>	<u>0.905</u>	<b>0.060</b>	<b>0.34</b>
Undistorted RFF	NeX [59]	<b>27.992</b>	<b>0.924</b>	<b>0.052</b>	0.13
	Ours (subdivision)	27.355	0.905	0.059	<b>0.34</b>
Shiny	NeX [59]	28.268	<b>0.940</b>	<b>0.036</b>	0.13
	Ours (subdivision)	<b>28.499</b>	0.931	0.038	<b>0.34</b>
Stanford	NeRF [30]	<u>37.559</u>	<u>0.979</u>	0.037	0.07
	NeX [59]	36.286	0.972	<u>0.033</u>	0.13
	X-Fields [2]	35.559	0.976	0.036	<b>20.00</b>
	Ours (no subdivision)	<b>38.054</b>	<b>0.982</b>	<b>0.020</b>	<u>9.14</u>

**Baseline Methods.** We compare our model to NeRF [30], AutoInt [22] (with 32 sections), and NeX [59] on the Real Forward-Facing dataset [30]; and additionally to NeX on the Shiny dataset [59]. As in [59], we perform evaluation on NeX before baking. More information about this is provided in the appendix. In Table 1, we include metrics for both the Real Forward-Facing and *Undistorted* Real Forward-Facing dataset, a variant of the original dataset that is supported by the NeX codebase. We did not include the following methods, although they are related, for various reasons: PlenOctrees [66], KiloNeRF [39], and NSVF [23] operate on bounded 360-degree scenes only; SNeRG [15] strictly underperforms NeRF on real forward-facing scenes.

On Stanford [58], we compare our model to NeRF, NeX, and X-Fields [2], which is a state of the art method in image-based light field interpolation. Kalantari et al. [16] would also be a relevant comparison, but X-Fields outperforms it, hence we did not include it in our evaluations.

**Subdivision.** For small-baseline light fields (Stanford [58]), *we do not use subdivision*, and use a single-evaluation per ray model with ray-space embedding. For sparse light fields (Real Forward Facing [30] and Shiny [59]) we place local light fields within each voxel in a  $32^3$  regular voxel grid that covers all of NDC space, with the exception of the CD and Lab sequences in the Shiny dataset, where we use a coarser  $4^3$  voxel grid.

**Implementation Details.** For our *feature-based embedding* (4), we set the embedding space dimension to  $N = 32$ . The embedded feature  $\mathbf{f}$  is  $\ell_2$ -normalized and then multiplied by  $\sqrt{32}$ , such that on average, each output feature has a square value close to one. For our *local affine transformation-based embedding* (5), we infer  $32 \times 4$  matrices and 32D bias vectors. We then normalize the matrices  $A$  with their Frobenius norm and multiply them by  $\sqrt{32} \cdot 4$ . We use a *tanh*-activation for our predicted bias vectors.

For positional encoding  $\gamma(\cdot)$  we use  $L = 8$  frequency bands for our subdivided model and  $L = 10$  frequency bands for our one-evaluation model. Higher frequencies are gradu-

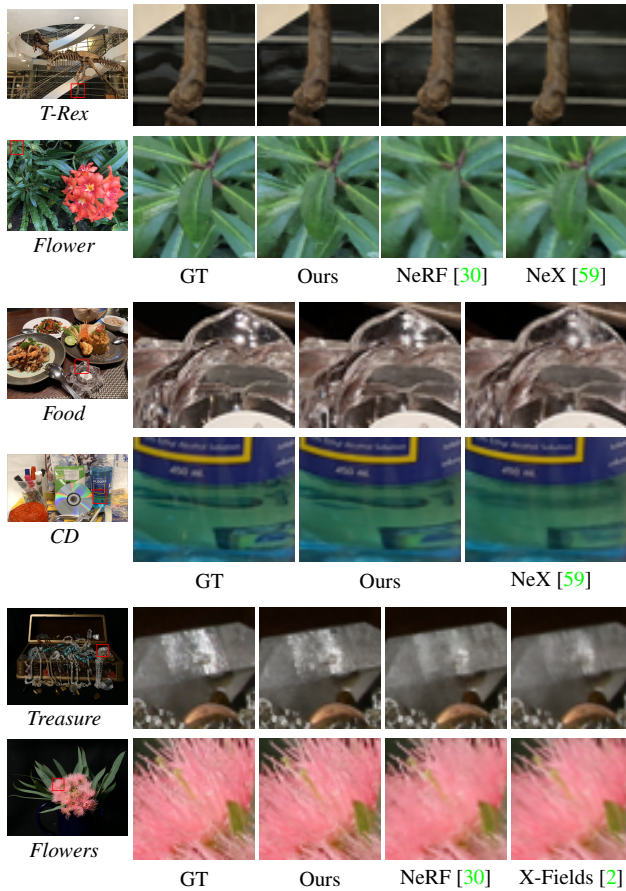


Figure 6. **Visual comparisons.** Close-ups reveal our representation leads to more faithful reconstruction of view-dependent effects, particularly distorted reflections and refractions (*Horns* and *T-Rex* from RFF; *Treasure* from Stanford; and *Food* and *CD* from Shiny) as well as sharper details (*Flower* from RFF and *Flowers* from Stanford). See our project website for more results.

ally eased in using windowed positional encoding, following Park et al. [35] — over 50k iterations with subdivision, or 80k without. For both the embedding and color models we use a 8-layer, 256 hidden unit MLP with one skip connection. We share a single network each for all local light fields for both the embedding  $E_\phi$  and color  $F_\theta$  in our subdivided representation, which are “indexed” (see section 5) with the voxel center points  $\mathbf{p}_i$ , i.e.,  $i \equiv \mathbf{p}_i$ .

We use a batch size of 4,096 for our subdivided model, and a batch size of 8,192 for our one-evaluation model during training and run all experiments on a single Tesla V100 GPU with 16 GB of memory. Please refer to the appendix and our project webpage for further details and results. We will release the source code and pre-trained models to facilitate future research.

## 7. Results

**Quantitative Evaluation.** Table 1 compares view synthesis quality using the metrics aggregated over held-out im-

ages in all scenes. For the Real Forward-Facing dataset, our method outperforms AutoInt in both quality ( $\sim 2$  dB) and speed ( $\sim 30\%$ ) while performing competitively with NeRF on PSNR and SSIM (within 0.5 dB in quality) and rendering more than  $\times 4.5$  times faster. We also outperform NeRF on LPIPS. Additionally, we are competitive with NeX on the *Undistorted* Real Forward-Facing dataset (within 0.65 dB in quality). For the Shiny dataset, we outperform NeX in terms of PSNR, while performing slightly worse on SSIM and LPIPS metrics. We do especially well on the *CD* and *Lab* scenes with challenging view dependence, outperforming NeX by  $> 3$  dB each.

On Stanford, our method outperforms NeRF, NeX, and X-Fields in all three metrics. Specifically, our model performs better on more complex scenes, e.g., *Treasure* and *Bulldozer*, with complicated geometry and high-frequency textures, while NeRF performs slightly better on simpler scenes like *Chess* and *Beans*. Our model renders a frame in about 0.11 seconds in this setup, which is faster than all other methods apart from X-Fields.

**Training Time.** For both RFF and Shiny, it takes about 20 hours to train our model at  $32^3$  resolution, while for the Stanford data, we train our model for about 10 hours. NeRF takes approximately 18 hours to train on all datasets, while NeX takes approximately 36 hours on 2 GPUs. For X-Fields, we double the default capacity of their model as well as the number of images used for interpolation; and train X-Fields for 8 hours rather than the maximum time of 172 minutes listed in their paper [2]. See the appendix for more details.

**Reconstruction of View Dependence.** Our approach reconstructs view dependence better than the baseline methods in *both* sparse and dense regimes (see Figures 1(left) and 6), due to the flexibility of the embedding network to predict different affine transforms for different regions of ray-space. This effectively produces *locally deforming* depth estimates, which can better model reflected and refracted content, for example. On the other hand, unlike our method, volumetric representations use a fixed global geometry, and are either late-conditioned on viewing direction [30], or use lower frequency basis functions for view dependence [59]. Thus, they struggle to represent high-frequency view dependent appearance that breaks multi-view color constraints.

In general, reproducing complex view dependence is a major challenge for volumetric representations, and is one of the primary contributions of our work. The difference is especially prominent in the Shiny dataset comparisons in Figure 1, as well as the results for *T-rex*, *Food*, *CD*, and *Treasure* in Figure 6.

**Memory Footprint.** Table 2 shows the model sizes. Our model has a similar memory footprint to NeRF. NeX explicitly represents albedo for 16 MPI layers, leading to larger memory consumption than pure coordinate-based neural representations. Overall, our method provides the best trade-off



Table 2. **Memory footprint** of our model compared to other methods in terms of the number of trainable weights in each network. For X-Fields [2] we include the memory footprint of the model, and of the training images which must be maintained in order to synthesize novel views.

Method	Ours	NeRF [30]	NeX [59]	AutoInt [22]	X-Fields [2]
Size (MB)	5.4	4.6	38.4	4.6	3.8 / 75

Table 3. **Ablations on the ray-embedding network** with subdivision (on Real Forward-Facing [30]) and without subdivision (on Stanford [58] light fields).

Subdivision	Method	Embedding	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Used (On RFF)	Ours	Not used (3)	26.696	0.888	0.079
		Feature (4)	26.922	0.891	0.073
		Local affine (5)	<u>27.454</u>	<u>0.905</u>	<b>0.060</b>
	NeRF [30]		<b>27.928</b>	<b>0.916</b>	<u>0.065</u>
Not used (On Stanford)	Ours	Not used (3)	25.111	0.841	0.098
		Feature (4)	37.120	<u>0.9792</u>	<u>0.030</u>
		Local affine (5)	<b>38.054</b>	<b>0.982</b>	<b>0.020</b>
	NeRF [30]		<u>37.559</u>	0.9790	0.037

between quality, rendering speed, and memory among the baselines we compared to; see the graphs in Figure 1(right).

**Ablations on Embedding Networks.** We study the effect of our proposed embedding networks via ablations. See Table 3 for quantitative evaluation with varying configurations.

With no embedding and only positional encoding (baseline), the network does not produce multi-view consistent view synthesis. The use of our feature-based embedding leads to a boost in quality over the baseline. In particular, it achieves our first objective of “memorization” (identifying rays that observe the same 3D point), but still struggles with view interpolation. Our local affine transformation embedding achieves better interpolation quality because it implicitly encourages color level sets to be locally linear. Thus, we learn a model that interpolates in a (locally) view consistent way for free.

**Ablations on Subdivision.** Table 4 and Figure 1(right) summarize the trade-off between the quality and rendering speed of our model on various resolutions of spatial subdivision, compared to NeRF and AutoInt. For each subdivision level, we use a regular  $N^3$  voxel grid, where  $N \in (4, 8, 16, 32)$ . As expected, our model archives better quality view synthesis as we employ more local light fields with finer spatial subdivision, trading off the rendering speed. The render-time increases linearly in  $N$ . Our method provides better quality and speed than AutoInt at *all* comparable subdivision levels.

**Limitations.** One limitation of our work is that we focus on light fields parameterized with two planes. Although common, this does not encompass  $360^\circ$  scenes. In the appendix, we show preliminary results for our method on  $360^\circ$  scenes using Plücker coordinates [47], which again outperforms Au-

Table 4. **Ablations on the subdivision** on sparse light fields (Real Forward-Facing [30]).

Method	Subdivision	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FPS $\uparrow$
Ours	$4^3$ grid	25.579	0.867	0.094	<b>2.82</b>
	$8^3$ grid	26.860	0.893	0.070	<u>1.42</u>
	$16^3$ grid	27.350	0.904	<u>0.062</u>	0.69
	$32^3$ grid	<u>27.454</u>	<u>0.905</u>	<b>0.060</b>	0.34
AutoInt [22]	8 sections	24.136	0.820	0.176	0.94
	16 sections	24.898	0.836	0.167	0.51
	32 sections	25.531	0.853	0.156	0.26
NeRF [30]		<b>27.928</b>	<b>0.916</b>	0.065	0.07

toInt. However, color level sets for Plücker coordinates are no longer affine. Thus, we believe that improving the design of embedding networks for different parameterizations could lead to a larger boost in performance.

Additionally, our representation requires subdivision for sparse light fields. This comes at the cost of both increased render time and training time. We hope to investigate other network designs or regularization schemes, which encourage the representation to be multi-view consistent without subdivision. Adaptive subdivision [23, 26] is another interesting direction for future work which could lead to better quality without sacrificing speed.

## 8. Conclusion

We present a novel ray-space embedding approach for learning *neural light fields*, which achieves state-of-the-art quality on small-baseline datasets. To better handle sparse input, we leverage spatial subdivision with a voxel grid of *local light fields*, which improves quality at the cost of increased render time. Our subdivided representation enables comparable performance to existing models, and achieves a better trade-off between quality, speed, and memory. Additionally, in both regimes, our method can handle complex view dependent effects that existing state-of-the-art volumetric scene representations do not faithfully reproduce [30, 59].

We believe that our approach is orthogonal to contemporary works that leverage hybrid or explicit grid-based representations for view synthesis [31, 50, 65], as well as approaches that use many *small* MLPs within a 3D volume [39]. As such, our approach can potentially be used to reduce the required grid resolution and number of sample points per ray for these methods — which may further improve rendering quality, speed, and memory overhead. We hope that our work can open up new avenues for view synthesis; and further, that its insights can be leveraged for other signal-representation problems (such as image/video representation) or high-dimensional interpolation problems (such as view-time or view-time-illumination interpolation [2]).



## References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. In *ECCV*, 2020. 3
- [2] Mojtaba Bemana, Karol Myszkowski, Hans-Peter Seidel, and Tobias Ritschel. X-fields: implicit neural view-, light- and time-image interpolation. *ACM Trans. Graph.*, 39(6):257:1–257:15, 2020. 3, 6, 7, 8
- [3] 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 Transactions on Graphics (TOG)*, 39(4):86–1, 2020. 3
- [4] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001. 2, 3
- [5] 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*, 2020. 3
- [6] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM Transactions on Graphics (TOG)*, 32(3):1–12, 2013. 3
- [7] Abe Davis, Marc Levoy, and Fredo Durand. Unstructured light fields. In *Eurographics*, 2012. 2
- [8] Brandon Yushan Feng and Amitabh Varshney. Signet: Efficient neural representation for light fields. In *ICCV*, 2021. 2, 3
- [9] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world’s imagery. In *CVPR*, 2016. 3
- [10] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, 2021. 3
- [11] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *ICCV*, 2021. 2
- [12] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996. 2
- [13] Peter Hedman, Suhil Alisan, Richard Szeliski, and Johannes Kopf. Casual 3d photography. *ACM Transactions on Graphics (TOG)*, 36(6):1–15, 2017. 3
- [14] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (TOG)*, 37(6):1–15, 2018. 3
- [15] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. *ICCV*, 2021. 2, 3, 6
- [16] Nima Khademi Kalantari, Ting-Chun Wang, and Ravi Ramamoorthi. Learning-based view synthesis for light field cameras. *ACM Transactions on Graphics (TOG)*, 35(6):1–10, 2016. 3, 6
- [17] Petr Kellnhofer, Lars Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. In *CVPR*, 2021. 2
- [18] Georgios Kopanas, Julien Philip, Thomas Leimkühler, and George Drettakis. Point-based neural rendering with per-view optimization. *Computer Graphics Forum*, 40(4):29–43, 2021. 3
- [19] Johannes Kopf, Kevin Matzen, Suhil Alisan, Ocean Quigley, Francis Ge, Yangming Chong, Josh Patterson, Jan-Michael Frahm, Shu Wu, Matthew Yu, et al. One shot 3d photography. *ACM Transactions on Graphics (TOG)*, 39(4):76–1, 2020. 3
- [20] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996. 2
- [21] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *CVPR*, 2021. 3
- [22] David B Lindell, Julien NP Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural volume rendering. In *CVPR*, 2021. 2, 3, 6, 8
- [23] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *NeurIPS*, 2020. 3, 6, 8
- [24] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Transactions on Graphics (TOG)*, 2021. 3
- [25] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. In *SIGGRAPH*, 2019. 3
- [26] Julien N. P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: Adaptive coordinate networks for neural scene representation. *ACM Transactions on Graphics (TOG)*, 40(4), 2021. 3, 8
- [27] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *CVPR*, 2021. 3
- [28] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019. 3
- [29] 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 Transactions on Graphics (TOG)*, 38(4):1–14, 2019. 3
- [30] 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, 5, 6, 7, 8
- [31] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022. 2, 8

- [32] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Joerg H. Mueller, Chakravarty R. Alla Chaitanya, Anton Kaplanyan, and Markus Steinberger. Donerf: Towards real-time rendering of compact neural radiance fields using depth oracle networks. *Comput. Graph. Forum*, 40(4):45–59, 2021. 2
- [33] Michael Oechsle, Michael Niemeyer, Christian Reiser, Lars Mescheder, Thilo Strauss, and Andreas Geiger. Learning implicit surface light fields. In *International Conference on 3D Vision (3DV)*, 2020. 3
- [34] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 3
- [35] 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 *ICCV*, 2021. 3, 7
- [36] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, 2021. 3
- [37] Eric Penner and Li Zhang. Soft 3d reconstruction for view synthesis. *ACM Transactions on Graphics (TOG)*, 36(6):1–11, 2017. 3
- [38] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. Derf: Decomposed radiance fields. In *CVPR*, 2021. 3
- [39] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *ICCV*, 2021. 2, 3, 6, 8
- [40] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *ECCV*, 2020. 3
- [41] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *CVPR*, 2021. 3
- [42] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *arXiv preprint arXiv:2110.06635*, 2021. 3
- [43] Lixin Shi, Haitham Hassanein, Abe Davis, Dina Katabi, and Fredo Durand. Light field reconstruction using sparsity in the continuous fourier domain. *ACM Transactions on Graphics (TOG)*, 34(1):1–13, 2014. 3
- [44] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3d photography using context-aware layered depth inpainting. In *CVPR*, 2020. 3
- [45] Harry Shum and Sing Bing Kang. Review of image-based rendering techniques. In *Visual Communications and Image Processing 2000*, volume 4067, pages 2–13, 2000. 3
- [46] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *NeurIPS*, 2020. 3
- [47] Vincent Sitzmann, Semon Rezchikov, William T Freeman, Joshua B Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. In *NeurIPS*, 2021. 2, 3, 8
- [48] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *CVPR*, 2019. 3
- [49] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *NeurIPS*, 2019. 3
- [50] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. *arXiv preprint arXiv:2111.11215*, 2021. 2, 8
- [51] Matthew Tancik, Pratul P Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *NeurIPS*, 2020. 4, 5
- [52] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. State of the art on neural rendering. *Computer Graphics Forum*, 39(2):701–727, 2020. 1, 3
- [53] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, Yifan Wang, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, Tomas Simon, Christian Theobalt, Matthias Niessner, Jonathan T. Barron, Gordon Wetzstein, Michael Zollhofer, and Vladislav Golyanik. Advances in neural rendering, 2021. 1
- [54] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhofer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *ICCV*, 2021. 3
- [55] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. In *CVPR*, 2020. 3
- [56] Shubham Tulsiani, Richard Tucker, and Noah Snavely. Layer-structured 3d scene inference via view synthesis. In *ECCV*, 2018. 3
- [57] Suren Vagharshakyan, Robert Bregovic, and Atanas Gotchev. Light field reconstruction using shearlet transform. *IEEE TPAMI*, 40(1):133–147, 2017. 3
- [58] Bennett Wilburn, Neel Joshi, Vaibhav Vaish, Eino-Ville Talvala, Emilio R. Antúnez, Adam Barth, Andrew Adams, Mark Horowitz, and Marc Levoy. High performance imaging using large camera arrays. *ACM Trans. Graph.*, 24(3):765–776, 2005. 2, 6, 8
- [59] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. NeX: Real-time view synthesis with neural basis expansion. In *CVPR*, 2021. 2, 3, 6, 7, 8
- [60] Gaochang Wu, Yebin Liu, Qionghai Dai, and Tianyou Chai. Learning sheared epi structure for light field reconstruction. *IEEE TIP*, 28(7):3261–3273, 2019. 3
- [61] Gaochang Wu, Mandan Zhao, Liangyong Wang, Qionghai Dai, Tianyou Chai, and Yebin Liu. Light field reconstruction using deep convolutional network on epi. In *CVPR*, 2017. 3
- [62] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint video. In *CVPR*, 2021. 3

- [63] Fanbo Xiang, Zexiang Xu, Milos Hasan, Yannick Hold-Geoffroy, Kalyan Sunkavalli, and Hao Su. Neutex: Neural texture mapping for volumetric neural rendering. In *CVPR*, 2021. 3
- [64] Henry Wing Fung Yeung, Junhui Hou, Jie Chen, Yuk Ying Chung, and Xiaoming Chen. Fast light field reconstruction with deep coarse-to-fine modeling of spatial-angular clues. In *ECCV*, 2018. 3
- [65] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. *arXiv preprint arXiv:2112.05131*, 2021. 2, 8
- [66] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. In *CVPR*, 2021. 2, 3, 6
- [67] Kai Zhang, Gernot Riegler, Noah Snaveley, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *CoRR*, abs/2010.07492, 2020. 2
- [68] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew Davison. In-place scene labelling and understanding with implicit scene representation. In *ICCV*, 2021. 3
- [69] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snaveley. Stereo magnification: learning view synthesis using multiplane images. *ACM Transactions on Graphics (TOG)*, 2018. 3