

NeuLF: Efficient Novel View Synthesis with Neural 4D Light Field

Zhong Li¹ Liangchen Song² Celong Liu¹ Junsong Yuan² Yi Xu¹
¹OPPO US Research Center, Palo Alto, California, USA
²University at Buffalo, State University of New York, Buffalo, New York, USA

<https://lizhong3232.github.io/neulf/>



Figure 1. Given a set of images captured in front of a scene using a handheld camera, **Neural 4D Light Field (NeuLF)** uses an implicit neural representation to learn the mapping from rays to color values. With the learned model, novel views can be synthesized by predicting the color of each ray. Moreover, with our framework, auto-refocus effect can be inherently generated by predicting depth for each ray.

Abstract

In this paper, we present an efficient and robust deep learning solution for novel view synthesis of complex scenes. In our approach, a 3D scene is represented as a light field, i.e., a set of rays, each of which has a corresponding color when reaching the image plane. For efficient novel view rendering, we adopt a two-plane parameterization of the light field, where each ray is characterized by a 4D parameter. We then formulate the light field as a 4D function that maps 4D coordinates to corresponding color values. We train a deep fully connected network to optimize this implicit function and memorize the 3D scene. Then, the scene-specific model is used to synthesize novel views. Different from previous light field approaches which require dense view sampling to reliably render novel views, our method can render novel views by sampling rays and querying the color for each ray from the network directly, thus enabling high-quality light field rendering with a sparser set of training images. Per-ray depth can be optionally predicted by the network, thus enabling applications such as auto refocus. Our novel view synthesis results are comparable to the state-of-the-arts, and even superior in some challenging scenes with refraction and reflection. We achieve this while maintaining an interactive frame rate and a small memory

footprint.

1. Introduction

Novel view synthesis has long been studied by the computer vision and computer graphics community. It has many applications in multimedia, AR/VR, gaming, etc. Traditional computer vision approaches such as multi-view stereo (MVS) and structure-from-motion (SfM) aim to build a geometric representation of the scene first. An alternative approach is image-based rendering [2, 8, 15], where no underlying geometric model or only a simple proxy is needed. These methods can achieve photorealistic rendering. However, a typical light field setup prefers a dense sampling of views around a scene. It thus limits practical use of such an approach.

With the recent advancement of neural rendering [37], photorealistic rendering with only a sparse set of inputs can be achieved. One approach is to use an explicit geometric representation of a scene reconstructed using a traditional computer vision pipeline and learning-based rendering. Object-specific or category-specific meshes or multi-plane images (MPI) [46] can be used as the representation. However, these explicit representations do not allow

a network to learn the optimal representation of the scene. To achieve this, volume-based representations can be used [19]. But they typically require a large amount of memory space, especially for complex scenes.

Memory-efficient implicit representations have gained interests from the research community. For example, surface-based implicit representations can achieve state-of-the-art results and can provide a high-quality reconstruction of the scene geometry [13]. However, surface-based representations face challenges when dealing with complex lighting and geometry, such as transparency, translucency, and thin geometric structures. More recently, volume-based implicit representation achieves remarkable rendering results (e.g., NeRF [24]) and inspires follow-up research. One drawback of NeRF, nevertheless, is the time complexity of rendering. NeRF, in its original form, needs to query the network multiple times per ray and accumulate color and density along the query ray, which prohibits real-time applications. Although there have been many efforts to accelerate NeRF, they typically require depth proxy to train or rely on additional storage to achieve faster rendering [10, 27, 43].

We propose an efficient novel view synthesis framework, which we call Neural 4D Light Field (NeuLF). We define a scene as an implicit function that maps 4D light field rays to corresponding color values directly. This function can be implemented as a Multilayer Perceptron (MLP) and can be learned using only a sparse set of calibrated images placed in front of the scene. This formulation allows the color of a camera ray to be learned directly by the network and does not require a time-consuming ray-marcher during rendering as in NeRF. Thus, NeuLF achieves 1000x speedup over NeRF during inference, while producing similar or even better rendering quality. Our light field setup limits the novel viewpoints to be on the same side of the cameras, e.g., front views only. Despite these constraints, we argue that for many applications such as teleconferencing, these are reasonable trade-offs to gain much faster inference speed with high-quality rendering and a small memory footprint.

Comparison with NeRF: Although our work is inspired by NeRF [24], there are some key distinctions. NeRF represents the continuous scene function as a 5D radiance field. Such a representation has redundancy, i.e., color along a ray is constant in free space. By restricting the novel viewpoints to be outside of the convex hull of the object, the 5D radiance field can be reduced to a light field in a lower dimension, e.g. 4D. Table 1, we summarize the differences between NeuLF and NeRF.

Moreover, NeuLF can also optionally estimate per ray depth by enforcing multi-view and depth consistency. Using depth, applications such as auto refocus can be enabled. We show state-of-the-art novel view synthesis results on benchmark datasets and our own captured data (Fig. 1). The

Table 1. The comparisons between our proposed NeuLF and NeRF [24].

	NeRF	NeuLF
Input	5D	4D
Output	radiance, density	color
Viewpoint range	360°	front views
Rendering method	raymarching	direct evaluation
Rendering speed	slow	fast
Memory consumption	small	small
High-quality rendering	yes	yes

comparisons with existing approaches also validate the efficiency and effectiveness of our proposed method. In summary, our contributions are:

- We proposed a fast and memory-efficient novel view synthesis pipeline, which solves the mapping from 4D rays to colors directly.
- Compared with the state-of-the-arts, our method is better than NeRF and NeX when the scene contains challenging refraction and reflection effects. In addition, our method only needs 25% of the original input on those challenge scenes to achieve a similar or even better quality.
- Application wise, the framework we proposed can optionally estimate depth per ray; thus enabling applications such as 3D reconstruction and auto refocus.

2. Related Work

Our work builds upon previous work in traditional image-based rendering and implicit-function-based neural scene representation. In the following sections, we will review these fields and beyond in detail.

Image-based Rendering For novel view synthesis, image-based rendering has been studied as an alternative to geometric methods. In the seminal work of light field rendering [15], a 5D radiance field is reduced to a 4D light field considering the radiance along a ray remains constant in free space. The ray set in a light field can be parameterized in different ways, among which two-plane parameterization is the most common one. Rendering novel views from the light field involves extracting corresponding 2D slices from the 4D light field. To achieve better view interpolation, approximate geometry can be used [2, 8, 40]. Visual effects of variable focus and variable depth-of-field can also be achieved using light field [11].

With the advancement of deep learning, a few learning-based methods have been proposed to improve the traditional light field. For example, LFGAN [4] can learn texture and geometry information from light field data sets and in turn predict a small light field from one RGB image. [22]

enables high-quality reconstruction of a light field by learning the geometric features hierarchically using a residual network. [41] integrates an anti-aliasing module in a network to reduce the artifacts in the reconstructed light field. Our method learns an implicit function of the light field and achieves high-quality reconstruction with a sparse input.

Neural Scene Representation Neural rendering is an emerging field. One of the most important applications of neural rendering is novel view synthesis. A comprehensive survey of the topic can be found in [37].

An explicit geometric model can be used as the representation of a scene. [28] creates a proxy geometry of the scene using structure from motion and multi-view stereo (MVS). Then, a recurrent encoder-decoder network is used to synthesize new views from nearby views. To improve blending on imperfect meshes from MVS, [9] uses predicted weights from a network to perform blending. A high-quality parameterized mesh of the human body [45] and category-specific mesh reconstruction [12] can also be used as the proxy. Recently, Multi-plane Image (MPI) [46] has gained popularity. [46] learns to predict MPIs from stereo images. The range of novel views is later improved by [35]. [6] uses learned gradient descent to generate an MPI from a set of sparse inputs. [23] uses an MPI representation for turning each sampled view into a local light field. NeX [39] represents each pixel of an MPI with a linear combination of basis functions and achieves state-of-the-art rendering results in real-time. MPI representation might typically lead to stack-of-cards artifacts. [32] trains a network to reconstruct both the geometry and appearance of a scene on a 3D grid. For dynamic scenes, Neural Volumes (NV) [19] uses an encoder-decoder network to convert input images into a 3D volume representation. [20] extends NV using a mixture of volumetric primitives to achieve better and faster rendering. While volume-based representations allow for learning the 3D structure, they require large memory space, especially for large scenes.

Implicit-function-based approaches provide memory-efficient alternatives to explicit representations, while still allowing learning the 3D structure of the scene. Implicit representations can be categorized as implicit surface-based and implicit volume-based approaches. SRN [33] maps 3D coordinates to a local feature embedding at these coordinates. Then, a trained ray-marcher and a pixel generator are used to render novel views. IDR [42] uses an implicit Signed Distance Function (SDF) to model an object on 3D surface reconstruction. Neural Lumigraph [13] provides even better rendering quality by utilizing a sinusoidal representation network (SIREN) to model the SDF.

Our work is inspired by NeRF [24], which uses a network to map continuous 5D coordinates (location and view direction) to volume density and view-dependent radiance. Recent works have extended NeRF to support novel illumina-

tion conditions [34], rendering from unstructured image collections from the internet [21], large-scale unbounded scenes [44], unknown camera parameters [38], anti-aliasing [1], deformable models [26], dynamic scenes [16], etc. A lot of effort has been put into speeding up rendering with NeRF. DOnERF [25] places samples around scene surfaces by predicting sample locations along each ray. However, transparent objects will pose issues and it requires ground-truth depth for training. FastNeRF [7] achieves 200fps by factoring NeRF into a position-dependent network and a view-dependent network. This allows efficient caching of network outputs during rendering. [43] trains a NeRF-SH network, which maps coordinates to spherical harmonic coefficients and pre-samples the NeRF-SH into a sparse voxel-based octree structure. These pre-sampling approaches sacrifice additional memory storage for speedups. NSVF [18] represents a scene using a set of NeRF-like implicit fields defined on voxels and uses a sparse octree to achieve 10x speedup over NeRF during rendering. KiloNeRF [27] decomposes a scene into a grid of voxels and uses a smaller NeRF for each voxel. Storage costs will increase when more networks are used. Using AutoInt [17], calculations of any definite integral can be done in two network evaluations; this achieves 10x acceleration, but rendering quality is decreased. Compared with these approaches, our method achieves 1000x speedup over NeRF by representing the scene with an implicit 4D light field without any additional pre-sampling or storage overhead.

Recently a concurrent work [31] propose to use a network to direct regress the mapping from the 6D Plücker coordinates to colors. It leverages meta-learning to enable view synthesis using a single image observation in ShapeNet dataset [3]. In contrast, we use 4D representation and conduct extensive experiments on real-world scenes. Another concurrent work [5] transforms 4D light field representation by leveraging Gegenbauer polynomials basis, and learning the mapping from this basis function to color, however, it requires dense narrow baseline input with planar camera arrangement.

3. Our Method

In Fig. 2, we illustrate the pipeline of our system. In the following sections, we will first briefly discuss the light field, followed by our NeuLF representation and the proposed loss functions. We will also discuss our training strategies.

3.1. 4D Light Field Representation

All possible light rays in a space can be described by a 5D plenoptic function. Since radiance along a ray is constant if viewed from outside of the convex hull of the scene, this 5D function can be reduced to a 4D light field [8, 15]. The most common parameterization is a two-plane model

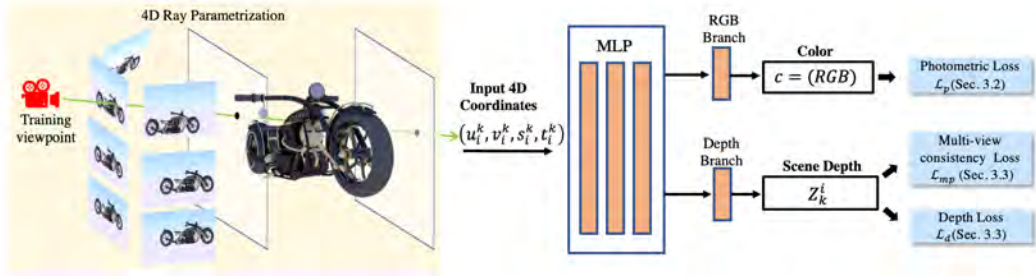


Figure 2. An overview of the Neural 4D Light Field (NeLF). For a set of sampled rays from training images, their 4D coordinates and the corresponding color values can be obtained. The input for NeLF is the 4D coordinate of a ray (query) and the output is its RGB color and scene depth. By optimizing the differences between the predicted colors and ground-truth colors, NeLF can faithfully learn the mapping between a 4D coordinate that characterizes the ray and its color. We also build a depth branch to let the network learn the per ray scene depth by self-supervised losses \mathcal{L}_{mp} and \mathcal{L}_d .

shown in Fig. 3. Each ray from the camera to the scene will intersect with the two planes. Thus, we can represent each ray using the coordinates of the intersections, (u, v) and (s, t) , or simply a 4D coordinate (u, v, s, t) . Using this representation, all rays from the object to one side of the two planes can be uniquely determined by a 4D coordinate.

Based on this representation, rendering a novel view can be done by querying all the rays from the center of projection to every pixel on the camera’s image plane. We denote them as $\{R_1, R_2, \dots, R_N\}$, where N is the total number of pixels. Then, for the i -th ray R_i , we can obtain its 4D coordinate (u_i, v_i, s_i, t_i) by computing its intersections with the two planes. If a function f maps the continuous 4D coordinates to color values, we can obtain the color of R_i by evaluating the function $f_c(u_i, v_i, s_i, t_i)$. In the next section, we will introduce Neural 4D Light Field (NeuLF) for reconstructing this mapping function f .

3.2. Neural 4D Light Field Reconstruction

We formulate the mapping function f_c as a Multilayer Perceptron (MLP). The input of this MLP is a 4D coordinate and the output is RGB color. As shown in Fig. 2. The




Figure 3. The 4D light field representation. Each ray is characterized by 4 parameters (u, v) and (s, t) , which uniquely locates the ray.

goal of the network is to learn the mapping function from training data.

Training Data: for a given scene, the training data comes from a set of captured images $\{I_1, I_2, \dots, I_M\}$, where M is the total number of images. Assuming the camera pose for each image is known or obtainable, for each image $I_k (k = 1, \dots, M)$, we can traverse its pixels and generate all corresponding rays $\{R_1^k, R_2^k, \dots, R_{N_k}^k\}$, where N_k is the total number of pixels in the k -th image. Based on the 4D light field representation, all 4D coordinates $\{(u_1^k, v_1^k, s_1^k, t_1^k), \dots, (u_{N_k}^k, v_{N_k}^k, s_{N_k}^k, t_{N_k}^k)\}$, ($k = 1, \dots, M$), can be obtained. On the other hand, the color for each pixel is known from the input images. To this end, we have constructed a collection of sample mappings from 4D coordinates to color values $(u_i^k, v_i^k, s_i^k, t_i^k) \rightarrow c_i^k$, ($k = 1 \dots M, i = 1 \dots N_k$), where c_i^k is the color of the i -th pixel on the k -th image. By feeding this training data to the MLP network, the parameters Θ can be learned by minimizing the following photometric loss \mathcal{L}_p :

$$\mathcal{L}_p = \sum_{k=1}^M \sum_{i=1}^{N_k} \|f_c(u_i^k, v_i^k, s_i^k, t_i^k | \Theta) - c_i^k\|_2 \quad (1)$$

In Fig. 2, we demonstrate an example of capturing images to train our neural 4D light field representation with a set of unstructured front-faced camera views. In this example, the cameras are placed on one side of two light slabs.

Rendering: Given a viewpoint \mathcal{V} , we can render a novel view $\mathcal{R}(\mathcal{V})$ by evaluating the learned mapping function f . With the camera pose and the desired rendering resolution $\{W^\mathcal{V}, H^\mathcal{V}\}$, we sample all rays $\{R_1^\mathcal{V}, R_2^\mathcal{V}, \dots, R_{N_\mathcal{V}}^\mathcal{V}\}$, where $N_\mathcal{V} = W^\mathcal{V} \times H^\mathcal{V}$ is the number of pixels to be rendered. We can further calculate the 4D coordinates $\{(u_i^\mathcal{V}, v_i^\mathcal{V}, s_i^\mathcal{V}, t_i^\mathcal{V})\}$ for each ray $R_i^\mathcal{V}$, ($i = 1 \dots N_\mathcal{V}$). We then formulate the rendering process \mathcal{R} as $N_\mathcal{V}$ evaluations of the mapping function f_c :




Figure 4. Given the 4D query ray $(u_i^k, v_i^k, s_i^k, t_i^k)$ (red line) and its predicted depth Z_i^k , we compute its corresponding rays (green line) to its nearest camera views C_i , which are used for self-supervision.

$$\mathcal{R}(\mathcal{V}) = \{f_c(u_i^v, v_i^v, s_i^v, t_i^v | \Theta), i = 1, \dots, N_V\}. \quad (2)$$

3.3. 4D Light Field Scene Depth Estimation

To simulate variable focus and variable depth-of-field using light field, previous work [11] dynamically reparameterizes the light field by manually moving the focal surface as a plane. To optimally and automatically select a focal plane given a pixel location, we aim to solve the focal surface manifold geometry that conforms to the scene geometry.

We let the same MLP predict per-ray depth as shown in Fig. 2. As in Fig. 4, for a 4D query ray $(u_i^k, v_i^k, s_i^k, t_i^k)$ from the camera C_k , the predicted scene depth is $Z_i^k = f_d(u_i^k, v_i^k, s_i^k, t_i^k | \Theta)$ where f_d is the network mapping function from 4D coordinates to scene depth. We compute ray-surface intersection (X_i^k, Y_i^k, Z_i^k) . Then, we self-supervise Z_i^k by applying multi-view consistency cues. To do this, we trace the rays from (X_i^k, Y_i^k, Z_i^k) back to C_k 's K -nearest data cameras $\{C_j\}, (j = 1, \dots, K)$ ($K=5$ in our experiments). Those rays intersect uv - st plane and can be parameterized as:

$$f_{uvst}(i, j, k, Z_i^k) = \{u_i^{kj}, v_i^{kj}, s_i^{kj}, t_i^{kj}\} \quad (3)$$

Ray-plane intersection is differentiable. We propose two loss functions \mathcal{L}_{mp} and \mathcal{L}_d to minimize the multi-view photometric error and depth differences respectively as follows:

$$\mathcal{L}_{mp} = \sum_{k=1}^M \sum_{i=1}^{N_k} \|f_c(u_i^k, v_i^k, s_i^k, t_i^k | \Theta) - \mathcal{C}_i^k(Z_i^k)\|_2 \quad (4)$$

$$\mathcal{C}_i^k(Z_i^k) = \sum_{j=1}^K \sum_{i=1}^{N_k} \omega(i, j) f_c(f_{uvst}(i, j, k, Z_i^k) | \Theta) \quad (5)$$

$$\mathcal{L}_d = \sum_{k=1}^M \sum_{i=1}^{N_k} \|f_d(u_i^k, v_i^k, s_i^k, t_i^k | \Theta) - \mathcal{D}_i^k(Z_i^k)\|_2 \quad (6)$$

$$\mathcal{D}_i^k(Z_i^k) = \sum_{j=1}^K \sum_{i=1}^{N_k} \omega(i, j) f_d(f_{uvst}(i, j, k, Z_i^k) | \Theta) \quad (7)$$

where $\omega(i, j)$ are the normalized weights $\omega(i, j) = \frac{1}{d_{ij}^2} / \sum_j \frac{1}{d_{ij}^2}$ with the Euclidean distance between the camera i and its neighbor camera j as d_{ij} . $\mathcal{C}_i^k(Z_i^k)$ and $\mathcal{D}_i^k(Z_i^k)$ are the supervise color and depth sum over weighted nearest data cameras with network output Z_i^k respectively. Training with both \mathcal{L}_{mp} and \mathcal{L}_d will encourage the MLP to learn the depth representation of the scene. Based on our experiments, only use one of above loss will lead to blurred depth and flickering artifacts. With ray-based depth, we can render auto-refocus effect given an arbitrary pixel location of the rendered novel view. The results are shown in Fig. 5 on our own captured data.




Figure 5. Depth estimation result and refocus effects. We show a novel view and its corresponding depth estimation. In addition, we show two auto-refocus results with focus point on the red cross (near) and on the green cross (far).

4. Experimental Results

We first discuss the implementation details of NeuLF. Then we perform quantitative and qualitative evaluations against state-of-the-art methods for novel view synthesis.

4.1. Implementation Details

We train the MLP on the following overall loss function:

$$\mathcal{L} = \mathcal{L}_p + \lambda_s \mathcal{L}_{mp} + \lambda_r \mathcal{L}_d, \quad (8)$$

where the weighting coefficients are $\lambda_s = 0.5$ and $\lambda_r = 0.1$. These parameters are fine-tuned by mixing the manual tuning and grid search tuning. This set of parameters works best for the scenes we tested. Grid search can be used to automatically find optimal parameters for a specific scene but will lead to a longer training time.

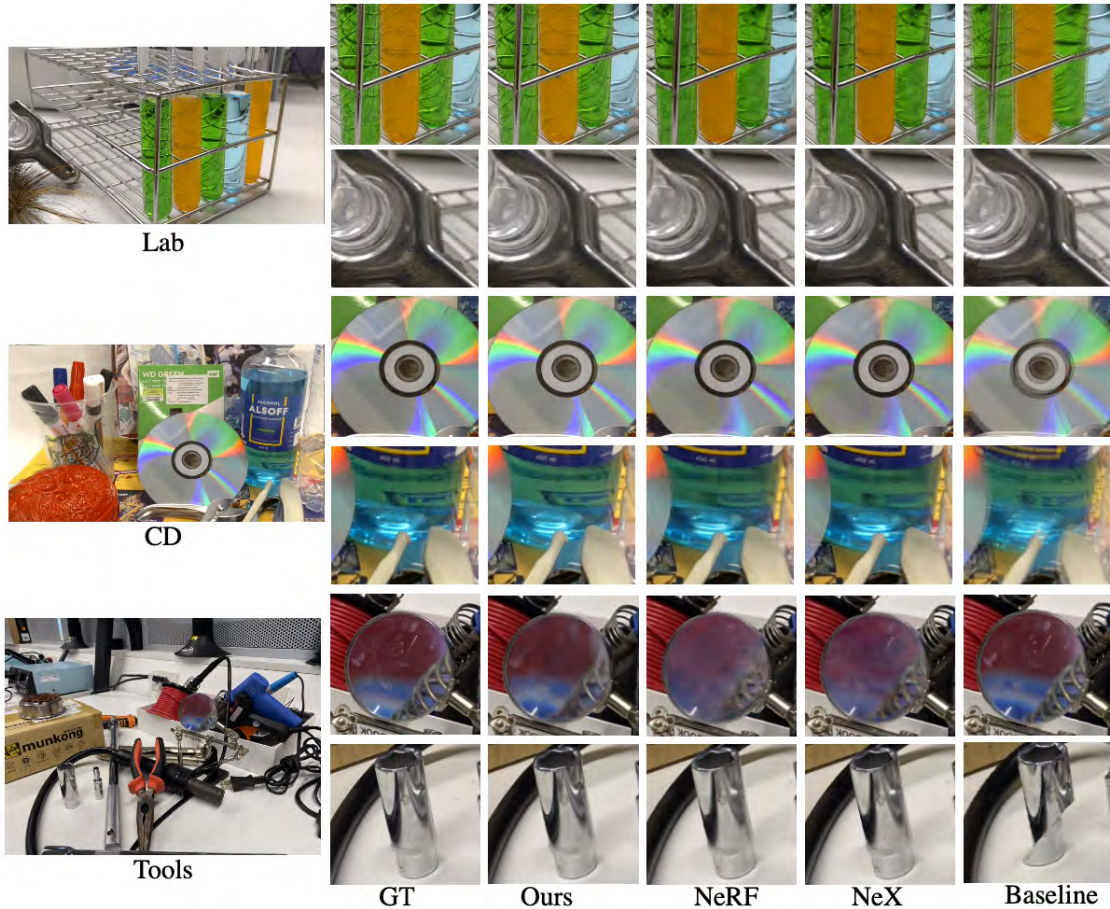


Figure 6. Qualitative Results on test views from shiny dataset. Our method captures more details on the reflection and refraction areas of the scenes.

Table 2. Average scores across test views for each scene in the shiny dataset.

	PSNR \uparrow				SSIM \uparrow				LPIPS \downarrow			
	Baseline	NeX	NeRF	Ours	Baseline	NeX	NeRF	Ours	Baseline	NeX	NeRF	Ours
Lab	21.69	30.43	29.60	31.95	0.693	0.949	0.936	0.951	0.261	0.146	0.182	0.097
CD	20.70	31.43	30.14	32.11	0.551	0.958	0.937	0.964	0.294	0.129	0.206	0.123
Giants	16.20	26.00	24.86	24.95	0.265	0.898	0.844	0.839	0.274	0.147	0.270	0.299
Tools	16.19	28.16	27.54	26.73	0.575	0.953	0.938	0.896	0.250	0.151	0.204	0.167
Food	14.57	23.68	23.32	22.61	0.297	0.832	0.796	0.776	0.341	0.203	0.308	0.322
Pasta	12.50	22.07	21.23	20.64	0.216	0.844	0.789	0.715	0.271	0.211	0.311	0.283
Seasoning	17.31	28.60	27.79	27.12	0.412	0.928	0.898	0.881	0.279	0.168	0.276	0.263
Crest	15.91	21.23	20.30	20.11	0.209	0.757	0.670	0.653	0.304	0.162	0.315	0.410

To extract camera rays from input photos, we calibrate the camera poses and intrinsic parameters using a structure-from-motion tool from COLMAP [29]. During training, we randomly select a batch of camera rays from the training set at each iteration. By passing them to the MLP to predict the color of each ray, we calculate \mathcal{L}_p and back-propagate the error.

The input 4D coordinate (u, v, s, t) (normalized to $[-1, 1]$) is passed through 20 fully-connected ReLU layers, each with 256 channels. We developed a structure that includes a skip connection that concatenates the input 4D coordinate to every 4 layers start with the fifth layer. An additional layer outputs 256-dimensional feature vector. This feature vector is split into the color branch and depth branch. Each branch is followed with an additional

fully-connected ReLU layer with 128 channels, and output 3 channel RGB radiance with sigmoid activation and 1 channel scene depth with sigmoid activation, respectively. For model training, we set the ray batch size in each iteration to 8192. We train the MLP for 1500 epochs using the Adam optimizer [14]. The initial learning rate is $1 \times 4^{10-4}$ and decays by 0.995 every epoch. To train the NeuLF on a scene with 30 input images with a resolution of 567×1008 , it takes 5 hours using 1 Nvidia RTX 3090 card. For testing on the same situation, rendering an image costs about 7ms while NeRF takes 51 seconds.

4.2. Comparison with State-of-the-Art Methods

In this section, we demonstrate the qualitative results of novel view synthesis and compare them with current top-performing approaches: NeRF [24], NeX [39] and the baseline light field rendering method [15]. We evaluate the models on the shiny dataset [39]. In this dataset, each scene is captured by a handheld smartphone in a forward-facing manner with a resolution 567×1008 . This is a challenging dataset that contains complex scene geometry and various challenging view-dependent effects. For example, refraction through the test tubes (filled with liquid) and magnifier, rainbow effect emits by a CD disk and sharp specular high-lights from silverware and thick glasses.

We hold out $\frac{1}{8}$ of each scene as the test set and the rest of them as the training set. The qualitative results are shown in Figure 6. The leftmost column shows our results on the test view of three challenging scenes (Lab, CD, and Tools). We have zoomed in on the part of the image areas for comparison with other methods. Our method is superior when a scene contains detailed refraction and reflection. Our result has less noise on glass sparkles. In the CD scene, our result shows more sharp and vivid detail on the rainbow, plastic cup reflections, and less noise on the liquid bottle than NeX and NeRF. In the Tools scene, although our result is not as sharp as the ground truth, it contains more overall details than NeX and NeRF and is able to recover metallic reflection with less noise than others. The baseline Light Field rendering (last column) exhibits good results when rays are sufficiently sampled (magnifier). However, its method exhibits aliasing and misalignment artifacts in the low ray sampling area (metallic, tube).

We report three metrics: PSNR (Peak Signal-to-Noise Ratio, higher is better), SSIM (Structural Similarity Index Measure, higher is better), and LPIPS (learned perceptual Image Patch Similarity, lower is better) to evaluate our test results. In Table 2, we report the three metrics for the 8 scenes in the shiny dataset. We use the NeX and NeRF scores originally report from the NeX paper. For each scene, we calculate the scores by averaging across the views in the test split. Our method produces the highest score across all three metrics on CD and Lab scenes

which contain challenging refraction and reflection. For the rest scenes, while NeX has the highest score by producing high-frequency details in the richly-textured area, we generate comparable scores with NeRF. Note that our rendering speed is $1000 \times$ than NeRF.

Table 3. Impacts from the number of inputs. We input different numbers of images to evaluate the performance of our model.

	method	#Images	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
CD 75%	Ours	230	31.81	0.959	0.126
CD 50%	Ours	153	31.41	0.953	0.145
CD 25%	Ours	77	30.16	0.948	0.170
CD 100%	NeRF	307	30.14	0.937	0.206
CD 100%	NeX	307	31.43	0.958	0.129
Lab 75%	Ours	227	31.87	0.949	0.097
Lab 50%	Ours	151	31.74	0.948	0.104
Lab 25%	Ours	76	30.61	0.939	0.116
Lab 100%	NeRF	303	29.60	0.936	0.182
Lab 100%	NeX	303	30.43	0.949	0.146

4.3. Study on Number of Inputs

To understand how the number of input views affect the novel view synthesis result, we train our model on fewer images. As shown in Figure 7, we use 75%, 50%, and 25% of the original data for the experiment. Although the input images are dramatically decreased, our method still generates high-quality results. Our results still retain the rainbow and background reflections, and the refraction details through the test tube are also well retained. In Table 3, note that even with less input images for CD and Lab scenes, our results are still comparable with or even better than NeRF and NeX with full number of inputs in the above challenging scenes.




Figure 7. Study on number of inputs. As shown in the figure, we input 75%, 50%, and 25% of the original images on the CD and Lab scenes.

4.4. Applications

As applications of NeuLF, we show results of depth estimation and automatic refocusing.

In Figure 8, we show an example of the depth estimation and refocusing effect on our own captured scene Tribe. Note the detailed depth of the house and statue are successfully recovered. With free-viewpoint scene depth, we can automatically select a focal plane given an image pixel location. We show two synthesized novel views rendered from our own captured scene. Then, we show the auto-refocus result given two positions on the image, one focuses on the near object (red cross), and another on the far object (green cross). This is enabled naturally by the dynamic 4D light field representation [11].




Figure 8. Depth estimation and auto-refocus result. We show the results of two novel views (first row) and the depth estimation (second row). In the third row we refocus on the near object, and in the fourth row we refocus on the far object.

4.5. Failure Cases

Our method is based on a 4D 2PP (two-plane parameterization) light field representation. Since each 3D world position corresponds to multiple discontinued 4D coordinates, NeuLF representation is difficult to learn. Therefore, we observe that our model cannot fully recover the high-frequency details in the scene as shown in Figure 9. The seaweed texture (first row) and object’s Hollow-out structure (second row) are over-smoothed. In addition, different exposure and lighting change across the frames can lead to flickering artifacts in the results. Recent works show that using a high-dimensional embedding [36], or using periodic activation functions [31] can help recover fine details. How-

ever, we found that the above method is causing over-fitting in the training views and a lack of accuracy in the test views on our NeuLF representation. Learning how to recover the fine details of the 2PP light field representation can be an interesting direction in the future.




Figure 9. Our failure cases. We show synthesized novel views and groundtruth. On the close-up views, we can see that our results lack fine details compared with groundtruth.

5. Conclusion

We propose a novel view synthesis framework called Neural 4D Light Field (NeuLF). Unlike NeRF, we represent a continuous scene using a 4D light field and train an MLP network to learn this mapping from input posed images. By limiting novel view synthesis to include only front views, NeuLF can achieve a comparable quality level as NeRF, but achieves a 1000x speedup. Moreover, because the speedup is enabled by modeling the color of light rays, NeuLF does not need additional storage for acceleration. To optionally output per-ray depth, we propose two loss terms: multi-view consistency loss and depth loss. This enables synthetic auto-refocus effect. We demonstrate state-of-the-art novel view synthesis results, especially for scenes with reflection and refraction. We also provide a study to show the effectiveness of our method with much fewer input images compared with NeRF and NeX.

6. Limitations and Future Work

There are several limitations to our approach. First, the novel viewpoints are limited to be on the one side of the two light slabs. In the future, we would like to extend the method to use more flexible 4D parameterizations such as multiple two planes, two cylindrical surfaces, or two spherical surfaces. By assuming the color is constant along a ray in free space, NeuLF cannot model rays that are blocked by the scene itself; therefore, novel viewpoints are always outside of the convex hull of the scene. This is an inherited limitation from light field.

Instead of using a 4D parameterization, lower-dimensional parameterization for specific applications can also be used. For example, in the work of concentric mosaic [30], by constraining camera motion to planar concen-

tric circles, all input image rays are indexed in three parameters. By adopting this parameterization, a more compact representation of the scene can be achieved, which potentially runs even faster than a 4D parameterization.

Free-viewpoint video can be a straightforward extension of NeuLF from static scenes to dynamic ones. In the future, we would like to explore the possibility of including time in the formulation following [16].

References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. *arXiv preprint arXiv:2103.13415*, 2021. [3](#)
- [2] 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*, SIGGRAPH '01, page 425–432, New York, NY, USA, 2001. Association for Computing Machinery. [1](#), [2](#)
- [3] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. [3](#)
- [4] Bin Chen, Lingyan Ruan, and Miu-Ling Lam. Lfgan: 4d light field synthesis from a single rgb image. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 16(1):1–20, 2020. [2](#)
- [5] Brandon Yushan Feng and Amitabh Varshney. Signet: Efficient neural representation for light fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14224–14233, 2021. [3](#)
- [6] John Flynn, Michael Broxton, Paul Debevec, Matthew Duvall, Graham Fyffe, Ryan Styles Overbeck, Noah Snavely, and Richard Tucker. Deepview: High-quality view synthesis by learned gradient descent. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [3](#)
- [7] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. *arXiv preprint arXiv:2103.10380*, 2021. [3](#)
- [8] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '96, page 43–54, New York, NY, USA, 1996. Association for Computing Machinery. [1](#), [2](#), [3](#)
- [9] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *37(6):257:1–257:15*, 2018. [3](#)
- [10] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. *arXiv preprint arXiv:2103.14645*, 2021. [2](#)
- [11] Aaron Isaksen, Leonard McMillan, and Steven J. Gortler. Dynamically reparameterized light fields. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, page 297–306, USA, 2000. ACM Press/Addison-Wesley Publishing Co. [2](#), [5](#), [8](#)
- [12] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 371–386, 2018. [3](#)
- [13] Petr Kellnhofer, Lars Jebe, Andrew Jones, Ryan Spicer, Kari Pulli, and Gordon Wetzstein. Neural lumigraph rendering. *arXiv preprint arXiv:2103.11571*, 2021. [2](#), [3](#)
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [7](#)
- [15] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42, 1996. [1](#), [2](#), [3](#), [7](#)
- [16] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, and Zhaoyang Lv. Neural 3d video synthesis, 2021. [3](#), [9](#)
- [17] David B. Lindell, Julien N. P. Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural volume rendering. In *Proc. CVPR*, 2021. [3](#)
- [18] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020. [3](#)
- [19] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 2019. [2](#), [3](#)
- [20] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering, 2021. [3](#)
- [21] Ricardo Martin-Brualla, Noha Radwan, Mehdi Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the wild: Neural radiance fields for unconstrained photo collections. <https://arxiv.org/abs/2008.02268>, 2020. [3](#)
- [22] Nan Meng, Xiaofei Wu, Jianzhuang Liu, and Edmund Lam. High-order residual network for light field super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11757–11764, 2020. [2](#)
- [23] 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](#)
- [24] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020. [2](#), [3](#), [7](#)

- [25] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Chakravarty R Alla Chaitanya, Anton Kaplanyan, and Markus Steinberger. Donerf: Towards real-time rendering of neural radiance fields using depth oracle networks. *arXiv preprint arXiv:2103.03231*, 2021. 3
- [26] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan Goldman, Steven Seitz, and Ricardo Martin-Brualla. Deformable neural radiance fields. <https://arxiv.org/abs/2011.12948>, 2020. 3
- [27] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps, 2021. 2, 3
- [28] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *European Conference on Computer Vision*, 2020. 3
- [29] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 6
- [30] Heung-Yeung Shum and Li-Wei He. Rendering with concentric mosaics. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 299–306, 1999. 8
- [31] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33, 2020. 3, 8
- [32] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2019. 3
- [33] Vincent Sitzmann, Michael Zollhoefer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems*, volume 32, 2019. 3
- [34] 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. *arXiv preprint arXiv:2012.03927*, 2020. 3
- [35] 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*, pages 175–184, 2019. 3
- [36] 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. *arXiv preprint arXiv:2006.10739*, 2020. 8
- [37] A. Tewari, O. Fried, J. Thies, V. Sitzmann, S. Lombardi, K. Sunkavalli, R. Martin-Brualla, T. Simon, J. Saragih, M. Nießner, R. Pandey, S. Fanello, G. Wetzstein, J.-Y. Zhu, C. Theobalt, M. Agrawala, E. Shechtman, D. B Goldman, and M. Zollhöfer. State of the Art on Neural Rendering. *Computer Graphics Forum (EG STAR 2020)*, 2020. 1, 3
- [38] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf —: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021. 3
- [39] Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3, 7
- [40] Daniel N Wood, Daniel I Azuma, Ken Aldinger, Brian Curless, Tom Duchamp, David H Salesin, and Werner Stuetzle. Surface light fields for 3d photography. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 287–296, 2000. 2
- [41] Gaochang Wu, Yebin Liu, Lu Fang, and Tianyou Chai. Revisiting light field rendering with deep anti-aliasing neural network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 3
- [42] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [43] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenotrees for real-time rendering of neural radiance fields. *arXiv preprint arXiv:2103.14024*, 2021. 2, 3
- [44] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 3
- [45] Xiuming Zhang, Sean Fanello, Yun-Ta Tsai, Tiancheng Sun, Tianfan Xue, Rohit Pandey, Sergio Orts-Escolano, Philip Davidson, Christoph Rhemann, Paul Debevec, et al. Neural light transport for relighting and view synthesis. *ACM Transactions on Graphics (TOG)*, 40(1):1–17, 2021. 3
- [46] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM Trans. Graph.*, 37(4), July 2018. 1, 3