

f NeRF: High Quality Radiance Fields from Practical Cameras

YI HUA, CHRISTOPH LASSNER, CARSTEN STOLL, and IAIN MATTHEWS, Epic Games, USA



Fig. 1. **Left:** a $5\text{cm} \times 8\text{cm} \times 6\text{cm}$ figurine captured with an $f/5.6$ aperture, (**center**) reconstructed with instant NeRF [Müller et al. 2022], and the proposed f NeRF (**right**). The reference image on the left was not used for reconstruction. While previous methods bake the defocus blur in to the radiance field, our method does not and reconstructs sharp details at all depths. It achieves this by successfully aggregating the radiance information from views close to the focus plane while simultaneously combining radiance values further away from the focus plane using a physically meaningful defocus model.

In recent years, the development of Neural Radiance Fields has enabled a previously unseen level of photo-realistic 3D reconstruction of scenes and objects from multi-view camera data. However, previous methods use an oversimplified pinhole camera model resulting in defocus blur being ‘baked’ into the reconstructed radiance field. We propose a modification to the ray casting that leverages the optics of lenses to enhance scene reconstruction in the presence of defocus blur. This allows us to improve the quality of radiance field reconstructions from the measurements of a practical camera with finite aperture. We show that the proposed model matches the defocus blur behavior of practical cameras more closely than pinhole models and other approximations of defocus blur models, particularly in the presence of partial occlusions. This allows us to achieve sharper reconstructions, improving the PSNR on validation of all-in-focus images, on both synthetic and real datasets, by up to 3 dB.

CCS Concepts: • **Computing methodologies** → **Rendering; Reconstruction**; Neural networks.

Additional Key Words and Phrases: Computational Photography, Multi-View & 3D, Differentiable Rendering, Neural Rendering

1 INTRODUCTION

Practical cameras are not pinhole cameras: they use a lens with a finite aperture to focus light on the sensor, resulting in higher light throughput and advantages such as faster exposure time and less vignetting. However, having a finite aperture inevitably results in defocus blur in the images. This occurs because real camera pixels integrate photons coming from different parts of the aperture, instead of a single infinitesimally small pinhole. Many modern radiance field reconstruction methods, such as Mildenhall et al. [2021], use such data with traditional pinhole camera models under the assumption of a pixel-to-ray correspondence. The resulting scene models have no choice but to ‘absorb’ the blur in erroneous material or density models.

Since multi-view data includes multiple observations of the same scene point, the same point may be seen in sharp focus in some frames and be blurry in others. If the forward imaging model closely matches the optics of a real camera lens, better scene reconstruction that preserves the sharp observation is possible. Many existing methods approximate the physics of light rays and follow the original Neural Radiance Field paper [Mildenhall et al. 2021], rendering pixel values by casting single rays from the pixel center through a pinhole. MipNeRF [Barron et al. 2021] and ZipNeRF [Barron et al. 2023] both propose sampling methods that are used to model a pixel’s integration of light from a cone; Wu et al. [2022] and Pidhorskyi et al. [2022] find 2D convolution approximations for defocus. However, none of these methods account correctly for partial occlusions (depth discontinuities) and they do not model the lens aperture and focus plane in a physically meaningful way. The ZipNeRF/MipNeRF cone aggregation strategies address sample aliasing, but incorrectly approximate partial occlusion within a pixel (for a visualization of the sample points on the cone surface, see Fig. 4). 2D convolution approximations are lacking information that is invisible in the optical center, but would be visible from other parts of the aperture. Methods for rendering physically-realistic camera measurements have been previously developed [Pharr et al. 2023]. Most prominently, one can use Monte-Carlo path tracing and trace multiple rays through the camera lens, refracting at the lens surface, to simulate the physical aggregation process of light at the sensor level. Presumably, previous radiance field reconstruction methods have not adapted it because the volumetric rendering process was already considered slow and additional samples would further slow down the reconstruction, making it intractable. Thanks to the recent progress in accelerating reconstruction and rendering for radiance fields [Božič et al. 2022; Kerbl et al. 2023; Müller et al. 2022], these



Fig. 2. Synthetic defocus on the MipNeRF360 ‘bicycle’ scene [Barron et al. 2022]. **Left** rendered with ZipNeRF [Barron et al. 2023] sampling cones modified to account for a larger aperture; **right** the proposed f NeRF. The modified ZipNeRF forward model doesn’t accurately model partial occlusions within a pixel, while f NeRF does and produces realistic blur and bokeh.

techniques become accessible—even more, we found that the practical runtime impact remains limited, presumably due to similarities in rays cast (for example, due to memory access patterns). Hence, we propose f NeRF, a method that leverages the physics of real lenses and their light aggregation process to enhance scene reconstruction.

Our contribution lies in adapting classical rendering techniques for volumetric rendering with finite aperture cameras. We show that this approach not only aligns more closely with optics, but also proves to be computationally tractable for the reconstruction of radiance fields. We also demonstrate that, using analytical gradient for aperture size, both aperture size and focus distance can be jointly optimized using gradient descent. This makes the proposed method easy to use in conjunction with many gradient-based 3D reconstruction techniques, such as neural radiance fields [Mildenhall et al. 2021; Müller et al. 2022], Gaussian splats [Kerbl et al. 2023] or implicit surfaces [Wang et al. 2021a].

In summary, we present a simple method for modeling lens defocus to significantly improve the sharpness of radiance field reconstruction from defocus blurred data. We demonstrate the validity of our approach on synthetic and real data with defocus blur.

2 RELATED WORK

2.1 Defocus blur in differentiable rendering

Several existing methods for gradient-based 3D reconstruction address the depth of field effect. [Mildenhall et al. 2021; Müller et al. 2022] render synthetic defocus on radiance fields, but do not model defocus in the reconstruction process. [Lee et al. 2024, 2023; Ma et al. 2022; Peng and Chellappa 2023] reconstruct radiance fields from blurred input images by optimizing the ray positions without a physical lens model. [Pidhorskyi et al. 2022; Wu et al. 2022] approximate the effect as a convolution with depth-dependent blur kernel, and show that jointly optimizing the aperture size and focus distance results in sharper reconstructions than no defocus blur modeling.

One practical consideration of the convolution-based approximations is that they require rendering of image patches proportional to blur size, instead of individual pixels, as well as careful handling of the convolution borders. More importantly, the convolution approximation is inaccurate along object boundaries and degrades reconstruction quality. We found that the physically-realistic approach to modeling defocus blur results in better reconstructions, while our implementation keeps the reconstruction time within a reasonable range (Fig. 6.). We show a comparison on defocused data from [Wu et al. 2022] in 5.2.2¹.

2.2 Camera modeling

Many papers show that better camera modeling improves the reconstruction of NeRFs from real data. [Lin et al. 2021; Park et al. 2023; Wang et al. 2021b] jointly optimize camera intrinsics and extrinsics while [Jeong et al. 2021; Xian et al. 2023] focus on the optimization of lens distortion. Since all these methods only cast a single ray from each pixel, they can all be easily extended with the proposed method for modeling finite aperture of practical cameras.

Inspired by the success of NeRF, recent camera modeling works are embracing the use of implicit networks with coordinate inputs. Some use a network to predict the blur at each pixel and use it to render the camera measurements [Lin et al. 2023], and demonstrate improvements on multiple reconstruction problems, including reconstructing all-in-focus image from image stacks [Huang et al. 2023; Wang et al. 2023]. These previous works capture a scene with different configurations of the camera, such as changing the exposure time or focus distance, while we take a camera with fixed configuration and focus on reconstructing a high quality radiance field of the scene. Additionally, these works are less physically grounded as they do not explicitly model ray casting from the aperture, so they can not optimize the aperture radius. In this paper we propose an efficient method for rendering the gradient of aperture radius.

2.3 Sampling volumes in NeRF

Multiple works [Barron et al. 2021, 2023] address the aliasing problem in rendering radiance fields. The need for anti-aliasing comes from the fact that each pixel measurement integrates photons across the pixel area as well as across the aperture. However, the cone modeling in MipNeRF and ZipNeRF, from which the radiance field is sampled, only accounts for sampling the pixel area and does not take into account the camera aperture.

To incorporate aperture sampling two modifications are needed. First, the apex of the cone should not be the pinhole lens position, but it should rather lie on the focus plane. Lens optics refract light passing the same point on the focus plane to the same pixel on the sensor plane. Second, the width of the cone should also be expanded so that it matches the aperture size (see Fig. 4). We evaluate the effectiveness of this modified version of ZipNeRF with correct focus point location and cone width matching the aperture, and compare its performance against the other models in Fig. 5, Table 1 and Table 2. The result shows ZipNeRF with aperture cone modifications results in sharper details in comparison to the original ZipNeRF.

¹We contacted Pidhorskyi et al. but were not able to get access to their data or code for establishing a comparison.

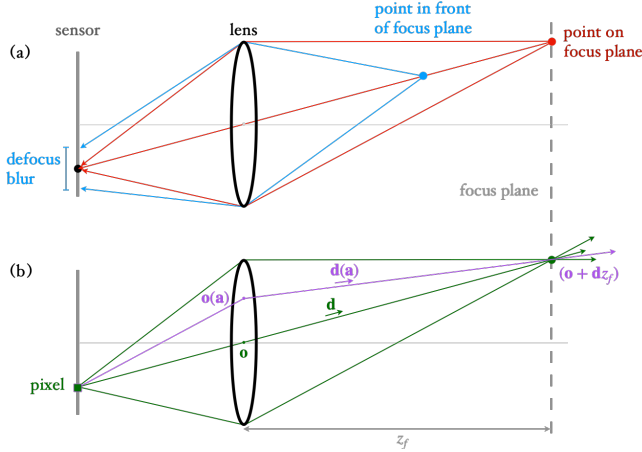


Fig. 3. Top: The red point on the focus plane results in a sharp image, while the blue point in front of the focus plane results in a blurry image, or bokeh. Bottom: to render color at pixel p , we draw samples a from the aperture and cast modified rays $(\mathbf{o}(a), \mathbf{d}(a))$ from it.

MipNeRF and ZipNeRF assume that all points in their sampled Gaussian share the same occlusion. This approximation allows for fast evaluation by allowing them to integrate in encoding space and therefore limiting the number of queries to the MLP. This is acceptable for approximating the sampling of the pixel area, however for larger apertures it is inaccurate for partial occlusions within a projection cone, which often occur at object boundaries. The difference is shown in Fig. 2. We show that with a more physically accurate model we can obtain better reconstructions for large apertures, in synthetic (Fig. 5) and real data (Fig. 8).

3 METHOD

In this section, we start by reviewing the mathematics of defocus blur. We then derive our method which extends volumetric rendering with aperture modeling, which allows us to reconstruct sharp radiance fields from defocus blurred measurements. For joint estimation of the defocus blur parameters with the radiance field, we also derive a method to efficiently render aperture radius gradient for circular apertures.

3.1 Defocus blur

Pinhole cameras have limited light throughput and suffer from heavy vignetting effects. In practice, cameras use a lens with a larger aperture to capture more light [Allen and Triantaphillidou 2011]. This allows either capturing with a shorter exposure time (important for moving subjects) or in dimly lit environments. However, a larger aperture creates defocus blur and hence leads to a loss of detail in areas further away from the focus plane [Potmesil and Chakravarty 1981]. Fig. 3(a) shows how the image of two points is formed on a camera with a finite aperture. For rays emitted by the red point on the focus plane, the lens refracts them so that all of them converge to the same point on the sensor, forming a sharp image of the point. However, for rays emitted by the blue point in front of the focus plane, the lens is unable to focus them and each ray will intersect

the sensor at a slightly different point, spreading out its image, an effect known as bokeh.

For a camera with aperture radius a_R , focal length f and focused at distance z_f , the radius of defocus blur, often known as circle of confusion, for scene point at depth z is,

$$a_R \frac{|z - z_f|}{z} \frac{f}{z_f - f}. \quad (1)$$

The defocus blur is the same for scene points at the same depth, thus previous methods [Pidhorskyi et al. 2022; Wu et al. 2022] approximate the defocus blur effect as 2D convolutions between a sharp image of the scene and a bokeh kernel. However, using 2D convolution approximations leads to inaccuracies at occlusion boundaries and requires discretizing a 3D scenes into depth layers. Volumetric rendering as used by many radiance field rendering methods offers us an opportunity to render higher fidelity defocus blur avoiding these problems.

3.2 Volumetric rendering with finite aperture

Most previous volumetric rendering methods model all cameras as pinhole cameras and only generate rays from the center of a pixel to the pinhole point. Specifically, the original NeRF [Mildenhall et al. 2021] expects color \tilde{C} from ray \mathbf{r} from a scene with volume density $\sigma(\mathbf{x})$ and emitted color \mathbf{c} ,

$$\tilde{C}(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t), \mathbf{d}) dt, \quad (2)$$

where $T(t) = \exp(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds)$, and t_n, t_f are respectively the near and far depth that bounds the scene. Modeling the lens aperture expects pixel color $C(\mathbf{p})$ integrated over many rays across different parts of the aperture \mathcal{A} . We follow classical rendering [Pharr et al. 2023] and sample rays from the aperture,

$$C(\mathbf{p}) = \frac{1}{\text{area}(\mathcal{A})} \int_{a \in \mathcal{A}} \tilde{C}(\mathbf{r}(a)) da, \quad (3)$$

where $\mathbf{r}(a)$ modifies a ray $\mathbf{r} = (\mathbf{o}, \mathbf{d})$ with origin \mathbf{o} at the pinhole and direction \mathbf{d} for finite aperture width. For focus distance of z_f , the modified ray $\mathbf{r}(a) = (\mathbf{o}(a), \mathbf{d}(a))$ has its origin on the aperture

$$\mathbf{o}(a) = \mathbf{a} \quad (4)$$

and direction $\mathbf{d}(a)$ pointing to the focus point,

$$\mathbf{d}(a) = \frac{\mathbf{o} + \mathbf{d}z_f - \mathbf{a}}{\|\mathbf{o} + \mathbf{d}z_f - \mathbf{a}\|_2}. \quad (5)$$

See Fig. 3 (b).

Our implementation samples points a from the aperture using a quasi-Monte Carlo approach using a Sobol Sequence [Sobol 1967]. This improves reconstruction performance, especially at low sample counts. If each pixel samples multiple rays, we add offsets to the points sampled during ray matching to reduce z -aliasing. For i -th ray among the n rays we sample for each pixel, we sample the point in ray section (t_{start}, t_{end}) at,

$$t = \frac{i}{n} t_{start} + \left(1 - \frac{i}{n}\right) t_{end}. \quad (6)$$

3.3 Optimizing aperture and focus depth

Highest quality reconstructions require accurate estimates of the camera parameters. Structure-from-motion packages together with EXIF data from recordings or manual recording of the parameters can provide great starting points. The dense, pixel-wise optimization of the entire scene provides ample opportunity to improve over these initializations and leads to better results [Park et al. 2023; Xian et al. 2023].

Our extended model has additional parameters for focus distance and aperture radius. The focus distance can naturally be optimized by using the auto-differentiation of the ray casting operation. The aperture radius, however, is a parameter for the sampling operation and its auto-differentiation gradient has high variance, due to the random sampling.

We follow [Pidhorskyi et al. 2022] and render the gradient of aperture radius from its analytic gradient. The gradient of pixel color $C(\mathbf{p})$ with respect to aperture radius a_R as,

$$\begin{aligned} \frac{dC(\mathbf{p})}{da_R} &= \frac{2}{a_R} C_{\text{ring}}(a_R) + 0 - \frac{2}{a_R} \int_{a_r=0}^{a_R} \frac{2a_r}{a_r^2} C_{\text{ring}}(a_r) da_r \quad (7) \\ &= \frac{2}{a_R} \left(C_{\text{ring}}(a_R) - C(\mathbf{p}) \right). \quad (8) \end{aligned}$$

Thus, we can render $C_{\text{ring}}(a_R)$ by sampling points on the boundary of the aperture and use Eq. (8) for computing the gradient of the aperture radius; this gives us stable convergence of the aperture radius, as shown in Table 3.

4 IMPLEMENTATION

We implement our framework using neural graphics primitives for fast reconstruction and rendering speeds and a compact volume representation [Müller et al. 2022]. Additionally, we use a dense volume acceleration structure for empty space skipping [Li et al. 2023]. Similar to NerfAcc, we use a sample-point target to guide ray batching and adjust the number of pixels sampled per batch dynamically. We use an ADAM optimizer [Kingma and Ba 2014] with an exponential learning rate reduction schedule with two steps to guide the optimization towards convergence. We observed that it is critical to average colors in linear space to create realistic bokeh effects, in particular in bright areas, and account for the logarithmic transfer function in sRGB encoded images during our reconstruction. We use the same batching, number of steps and learning rate schedule across all methods for a fair comparison. The experiments are run on NVIDIA A100 GPUs with 80GB memory unless otherwise specified.

5 EXPERIMENTS

5.1 Synthetic experiments

5.1.1 Experiment setup. We validate our method on the synthetic data rendered with large aperture; see Fig. 5. We render the blender scenes from [Mildenhall et al. 2021] in BlenderNeRF² [Raafat 2023] with $f/5$ and focus distance 3.5m for 100 images captured from a hemisphere with radius 4m. We render the same images sharp and without defocus for testing.

²Blender computes f numbers differently from photography conventions; to reproduce the rendering use $f/0.2$ and focal length 50mm camera in Blender.

scene	NGP	ZN	ZN + A	LN 6	LN 32
lego	0.8801	0.8806	0.8800	0.9214	0.9360
chair	0.9023	0.9023	0.9303	0.9261	0.9419
figus	0.9300	0.9395	0.9361	0.9645	0.9657
drums	0.8826	0.8838	0.6085	0.9037	0.9033
mic	0.9157	0.9157	0.9299	0.9325	0.9377
ship	0.7352	0.7362	0.7017	0.7515	0.7467
hotdog	0.9373	0.9377	0.9361	0.9549	0.9572

Table 1. Reconstruction SSIM from synthetic defocused measurements with different methods.

To isolate the effect of addressing defocus, we only use the reconstruction loss between measured pixel color C and rendered pixel color $C(\mathbf{p})$,

$$\ell_{\text{recon}} = \text{smooth } \ell_1(C(\mathbf{p}) - C), \quad (9)$$

where smooth ℓ_1 loss equals $\|x\|_2^2$ for $|x| \leq 1$ and $|x|$ otherwise; we use it for robust optimization. We use the ground truth camera intrinsics, extrinsics, aperture radius, and focus distance unless otherwise specified. All experiments are run with learning rate 1e-2 for a fixed number of 2e4 steps. We use base resolution of 16 with 16 levels and a maximum resolution of 4096 for the NGP volume, backed by a hashmap of size 2^{19} . The target number of point samples per batch is 262,144 for all methods.

5.1.2 Comparison against ZipNeRF and ZipNeRF modified for aperture. We implemented the multisampling and downweighting from ZipNeRF (ZN), as well as a modified version where we shift the sampling cone apex to the focus plane and expand the cone to match the aperture (ZN+A) for comparison with the proposed f NeRF method. The modification in the cone shape is shown in the sampling locations in Fig. 4. We report the SSIM and PSNR metrics on the sharp testing images in Table 3. An Instant NGP (NGP) reconstruction with pinhole camera model does not model lens aperture and bakes in the defocus blur in the scene, as shown on second column. The ZipNeRF (ZN) reconstruction models a cone defined by the pixel, not the lens aperture, and fails on large aperture scenarios as shown. ZipNeRF with the sampling cone modified for aperture (ZN+A) produces sharper results, but the model mismatch in partial occlusions tends to produce over-smoothed rendering and an over-sharpened reconstruction. ZN+A fails to reconstruct the drum scene due to the large amount of partial occlusions and specularities in this scene. Our proposed method with aperture modeling significantly improves the sharpness of reconstruction, with 32 rays (LN-32) showing more details than 6 rays (LN-6) on most scenes, as shown on the last two columns.

5.1.3 Number of rays per pixel. To understand how the reconstruction quality and runtime scale with the number of rays we sample for each pixel, we reconstruct the synthetic Lego scene with the proposed f NeRF method with different numbers of rays per pixel, but with a constant number of pixels per batch (1024) Fig. 6. This experiment is run on a NVIDIA Quadro RTX 8000. The reconstruction quality saturates near 16 rays per pixel. We also see that the runtime does not increase significantly below 8 rays per pixel; we posit that the lower increase in runtime originates from synergy




Fig. 4. Schematic of sampling locations of different methods in 2D. (a) blue points show ZipNeRF samples on a small cone with apex on the pinhole location; (b) yellow points show ZipNeRF modified for aperture by moving the apex to focus plane and expanding the cone to match aperture; (c) purple points show f NeRF samples on 6 rays, drawn at random, passing through the aperture. ZipNeRF sample solely on a cone surface whereas the proposed method casts rays within the cone volume. We show actual sampling locations in 3D in the supplementary video.

scene	NGP	ZN	ZN + A	LN 6	LN 32
lego	25.75	25.84	26.02	28.26	29.00
chair	27.25	27.24	27.84	28.9	29.64
figus	27.28	27.65	26.32	30.65	30.49
drums	22.54	22.46	14.53	23.93	23.43
mic	24.18	24.17	26.10	25.84	26.07
ship	22.64	22.68	21.46	23.51	23.17
hotdog	29.63	29.69	30.06	31.93	32.08

Table 2. Reconstruction PSNR from synthetic defocused measurements with different methods.

effects of the many rays in each batch being closely related in terms of memory accesses and computation.

5.1.4 Optimizing aperture radius and focus distance. To evaluate the convergence of the aperture radius and focus distance parameters, we reconstruct the synthetic Lego scene with mismatched initial parameters (80%, 100%, 120%). As shown in Tab. 3 mismatched defocus parameters result in degraded reconstruction quality, in some cases leading to complete failures to reconstruct the scene. Allowing for joint optimization of the defocus parameters with the scene leads to comparable quality with correctly optimized defocus parameters up to a range of 20% of the aperture and focus depth initialization. This means that in practice, the proposed method can be expected to work reliably even for ballpark estimates of these parameters.

aperture init.	focus init.	scene only	+ opt. defocus params.
80%	100%	28.73 0.7606	29.19 0.9295
<u>100%</u>	<u>100%</u>	28.98 0.9294	29.07 0.9311
120%	100%	27.72 0.8364	29.13 0.9276
100%	80%	18.55 0.7385	24.14 0.9341
<u>100%</u>	<u>100%</u>	29.04 0.9368	29.00 0.9366
100%	120%	22.86 0.8199	28.97 0.9338

Table 3. Reconstruction PSNR|SSIM from different aperture radius and focus distance initializations for the synthetic lego scene, for scene only and joint optimization with defocus parameters. The aperture and focus initialization is perturbed as specified in the leftmost columns. When optimizing the defocus parameters we manage to maintain near perfect reconstruction quality. Note that there is a slight variance due to the stochastic optimization in our method comparing the two lines initialized with ground truth.

5.2 Experiments on Recorded Data

5.2.1 Experiment setup. To evaluate reconstruction on recorded data, we optimize camera poses using CamP [Park et al. 2023] (using SE3, sampling 128 points in the reconstruction volume, with diagonal weight $\lambda = 1e-1$, and identity weight $\mu = 1e-8$) and the distortion loss from [Barron et al. 2022] ($\lambda_{dist} = 1.5938e-3$).

5.2.2 DoF NeRF experiments. We validate our method on data from DoF-NeRF [Wu et al. 2022], which captures scenes with different focus distances and different apertures ($f/4$ and $f/16$). We show comparisons with DoF-NeRF in Fig. 7. The input consists of images captured with 2 focus distances with the same camera pose, each consisting of 22 images. We run our instant NGP implementation with pinhole camera model as a baseline. For DoF-NeRF, since their code is unavailable, we use the PSNR number and result presented in their paper. We follow the configuration described in the DoF-NeRF paper and reconstruct with f NeRF at the same resolution, with input dimension of 497×322 , and jointly optimize aperture radius and focus distances. However, unlike DoF-NeRF which estimates the per-frame defocus

5.2.3 Mip-NeRF360 experiments. We also evaluate our method on the Mip-NeRF360 dataset. While the apertures are fairly small, we can still improve and produce sharper details in the reconstruction in Fig. 8. Since the captures are centered around an object, most background points were similarly defocused across all frames. The quantitative validation metrics are similar to baseline due to the validation set sharing the same defocus blur as input. Since each scene contains more than a hundred images at high resolution, we find a balance between pixels per batch and rays per pixel, by starting at randomly sampling 1 ray per pixel and doubling it for each epoch. The kitchen scene is reconstructed with $5e-4$ aperture radius and 1.09 focus distance, and the bicycle scene with $1.5e-4$ aperture radius and 0.15 focus distance. We use three sets of hyperparameters across all experiments in the paper: we adjust them between synthetic data, indoor and outdoor spaces. To provide a rough overview of their ranges: the max resolution of the NGP is between 4096 and 467830, backed by a hashmap of size $2^{19} - 2^{21}$. The overall learning rate for the ADAM optimizer is chosen between 0.01 and 0.0001, pose between $1e-5$ and $1e-4$, and intrinsics between $4e-6$ and $8e-6$. We also regularize the deviation from initial




Fig. 5. Synthetic data results. From left to right: input frame closest to test viewpoint, reconstruction from iNGP, ZipNeRF, ZipNeRF modified for aperture, the proposed method LensNeRF with 6 rays per pixel, LensNeRF with 32 rays per pixel, and all-in-focus ground truth. The closest input views demonstrate notably blurry regions in important areas of the image, and our reconstructed model is not able to leverage this viewpoint for the reconstruction of this viewing angle. However, other viewpoints were sufficient to reconstruct the areas in question with high fidelity and sharpness and the model is not negatively affected by the reconstruction loss thanks to our accurate depth-of-field model.

camera pose, with a rotation loss using the ℓ_2 distances between initial and optimized rotation matrices weighted by $\lambda_{\text{rot}}=9.59\text{e-}3$, and a translation loss using the ℓ_2 distance between initial and optimized camera locations weighted by $\lambda_{\text{trans}}=9.09\text{e-}3$. The learning rate schedule is being applied with two reduction steps with a size between 0.2 and 0.4. Other parameters include weight decay ($2\text{e-}6$ to $9\text{e-}6$) and gradient scaling ($600\text{x-}6000\text{x}$) for better numerical stability for compatibility with FP16 tensor operations.

5.2.4 New data. We image a miniature figurine of size $5\text{cm} \times 8\text{cm} \times 6\text{cm}$ with a Canon EOS M100 with $f/5.6$. Since the subject is small,

we capture it with the camera handheld close to it, resulting in a close focus distance and a shallow depth of field. We captured 96 images spread around the upper hemisphere. We reconstruct using 88 frames of 10 times downsampled images, at 600×400 resolution. One of the practical problems of reconstructing from defocus blurred data is that the initial camera poses from structure from motion are often inaccurate. This is not a problem if the images are from a pre-calibrated multi-camera capture system. DoF-NeRF solves this problem by capturing with both large and small aperture from the same positions and using the poses estimated only from small aperture photos. We do not capture any photos at small apertures




Fig. 6. Reconstruction quality and runtime v.s. number of rays per pixel on synthetic lego data. The reconstruction quality saturates near 16 rays per pixel; the runtime does not decrease significantly below 8 rays per pixel.




Fig. 7. DoF-NeRF dataset reconstructions from ($f/4$) captures. The numbers are PSNRs evaluated against $f/16$ images at 2 unseen poses.

and found that we can get reasonable pose estimates using Agisoft Metashape³ using downsampled data. This results in high quality reconstructions that we can segment in 3D. You can find a held out image, as well as reconstruction renders in Fig. 1.

6 CONCLUSION

In this paper, we improved the camera modeling in radiance field methods and pushed them closer to the realities of practical camera systems. The proposed method, LensNeRF, can be easily integrated into most radiance field reconstruction systems. It casts multiple rays from the camera aperture to render physically realistic defocus blur. We show that LensNeRF reconstructs sharp radiance fields and is computationally tractable. By incorporating the optics of a thin lens into our modeling, we are able to render the analytic gradient of aperture radius. We are successful at jointly estimating the aperture radius and focus distance, two parameters that control the amount of defocus, along with the scene. We have critically evaluated the current state of art volume sampling approach, ZipNeRF, and illustrated the effects of approximations included in that approach. We introduce an aperture sampling extension to ZipNeRF to more reasonably compare with our LensNeRF approach in scenarios involving partial occlusions. Our experiments on synthetic and real defocused camera inputs show up to 3dB PSNR improvement when evaluated on all-in-focus novel views.

³<https://www.agisoft.com>

While this work represents a stride toward more accurate modeling of practical camera systems, it is important to acknowledge its limitations. One of the key simplification is the use of the thin lens model, which, although effective, does not encapsulate the complexity of modern camera lenses with multiple elements. Some unique characteristics of cameras, such as chromatic aberration, lens flare, and bokeh variation cross different regions of the lens, are not represented in our model. Further research could focus on integrating more expressive camera models that capture these nuances while maintaining rendering efficiency. These advancements will push the fidelity of reconstructions from practical camera inputs.

REFERENCES

- Elizabeth Allen and Sophie Triantaphyllidou (Eds.). 2011. *The Manual of Photography* (tenth ed.). Focal Press.
- Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. 2021. Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, Montreal, QC, Canada, 5835–5844. <https://doi.org/10.1109/ICCV48922.2021.00580>
- Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. 2022. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. *CVPR* (2022).
- Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. 2023. Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields. *ICCV* (2023).
- Aljaž Božič, Denis Gladkov, Luke Doukakis, and Christoph Lassner. 2022. Neural Assets: Volumetric Object Capture and Rendering for Interactive Environments. arXiv:2212.06125 [cs.CV]
- Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, and Qing Wang. 2023. Inverting the Imaging Process by Learning an Implicit Camera Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 21456–21465.
- Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Animashree Anandkumar, Minsu Cho, and Jaesik Park. 2021. Self-Calibrating Neural Radiance Fields. In *ICCV*.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics* 42, 4 (July 2023). <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
- Diederik Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations* (2014).
- Byeonghyeon Lee, Howoong Lee, Usman Ali, and Eunbyung Park. 2024. Sharp-NeRF: Grid-based Fast Deblurring Neural Radiance Fields Using Sharpness Prior. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 3709–3718.
- Dogyoon Lee, Minhyeok Lee, Chajin Shin, and Sangyoun Lee. 2023. DP-NeRF: Deblurred Neural Radiance Field With Physical Scene Priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 12386–12396.
- Ruilong Li, Hang Gao, Matthew Tancik, and Angjoo Kanazawa. 2023. NerfAcc: Efficient Sampling Accelerates NeRFs. *arXiv preprint arXiv:2305.04966* (2023).
- Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. 2021. BARF: Bundle-Adjusting Neural Radiance Fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 5741–5751.
- Esther YH Lin, Zhecheng Wang, Rebecca Lin, Daniel Miao, Florian Kainz, Jiawen Chen, Xuaner Cecilia Zhang, David B Lindell, and Kiriakos N Kutulakos. 2023. Learning lens blur fields. *arXiv preprint arXiv:2310.11535* (2023).
- Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. 2022. Deblur-nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12861–12870.
- Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2021), 99–106.
- Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans. Graph.* 41, 4, Article 102 (July 2022), 15 pages. <https://doi.org/10.1145/3528223.3530127>
- Keunhong Park, Philipp Henzler, Ben Mildenhall, Jonathan T. Barron, and Ricardo Martin-Brualla. 2023. CamP: Camera Preconditioning for Neural Radiance Fields. *ACM Trans. Graph.* (2023).
- Cheng Peng and Rama Chellappa. 2023. PDRF: progressively deblurring radiance field for fast scene reconstruction from blurry images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 2029–2037.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically based rendering: From theory to implementation*. MIT Press.
- Stanislav Pidhorskyi, Timur Bagautdinov, Shugao Ma, Jason Saragih, Gabriel Schwartz, Yaser Sheikh, and Tomas Simon. 2022. Depth of Field Aware Differentiable Rendering.

- ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–18.
- Michael Potmesil and Indranil Chakravarty. 1981. Synthetic Image Generation with and Aperture Camera Model. *Computer Graphics* 15, 3 (1981).
- Maxime Raafat. 2023. *BlenderNeRF*. <https://doi.org/10.5281/zenodo.7926211>
- I Sobol. 1967. The distribution of points in a cube and the accurate evaluation of integrals (in Russian) *Zh. Vychisl. Mat. i Mater. Phys* 7 (1967), 784–802.
- Chao Wang, Ana Serrano, Xingang Pan, Krzysztof Wolski, Bin Chen, Hans-Peter Seidel, Christian Theobalt, Karol Myszkowski, and Thomas Leimkühler. 2023. A Neural Implicit Representation for the Image Stack: Depth, All in Focus, and High Dynamic Range. *ACM Transactions on Graphics* 42, 6 (2023). <https://doi.org/10.1145/3618367>
- Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. 2021a. NeuS: Learning Neural Implicit Surfaces by Volume Rendering for Multi-view Reconstruction. *NeurIPS* (2021).
- Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. 2021b. NeRF—: Neural Radiance Fields Without Known Camera Parameters. *arXiv preprint arXiv:2102.07064* (2021).
- Zijin Wu, Xingyi Li, Juewen Peng, Hao Lu, Zhiguo Cao, and Weicai Zhong. 2022. Dofnerf: Depth-of-field meets neural radiance fields. In *Proceedings of the 30th ACM International Conference on Multimedia*. 1718–1729.
- Wenqi Xian, Aljaž Božič, Noah Snavely, and Christoph Lassner. 2023. Neural Lens Modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.




Fig. 8. Results on MipNerf360 [Barron et al. 2022] data. Both NGP and ZipNeRF bake defocus blur into their reconstructions, resulting in smoothed out details. ZN+A generates over-sharpened details and introduces artifacts near occlusion boundaries. fNeRF reconstructs better details without artifacts. Please refer to our supplementary results for a better visualization of the reconstructions.