

Cinematic Gaussians: Real-Time HDR Radiance Fields with Depth of Field

Chao Wang¹, Krzysztof Wolski¹, Bernhard Kerbl^{2,3}, Ana Serrano⁴, Mojtaba Bermana¹,
 Hans-Peter Seidel¹, Karol Myszkowski¹, and Thomas Leimkühler¹

¹Max-Planck-Institut für Informatik, Germany

²Technische Universität Wien, Austria

³Carnegie Mellon University, USA

⁴Universidad de Zaragoza, I3A, Spain



Figure 1: Based on sparse images captured from various viewpoints and with different focus distances, apertures, and exposure times (left), our method generates an HDR 3D Gaussian representation that can be rendered with arbitrary depth of field. Our representation supports rendering images in a six-dimensional space of camera and lens parameters (right, five dimensions are shown).

Abstract

Radiance field methods represent the state of the art in reconstructing complex scenes from multi-view photos. However, these reconstructions often suffer from one or both of the following limitations: First, they typically represent scenes in low dynamic range (LDR), which restricts their use to evenly lit environments and hinders immersive viewing experiences. Secondly, their reliance on a pinhole camera model, assuming all scene elements are in focus in the input images, presents practical challenges and complicates refocusing during novel-view synthesis. Addressing these limitations, we present a lightweight method based on 3D Gaussian Splatting that utilizes multi-view LDR images of a scene with varying exposure times, apertures, and focus distances as input to reconstruct a high-dynamic-range (HDR) radiance field. By incorporating analytical convolutions of Gaussians based on a thin-lens camera model as well as a tonemapping module, our reconstructions enable the rendering of HDR content with flexible refocusing capabilities. We demonstrate that our combined treatment of HDR and depth of field facilitates real-time cinematic rendering, outperforming the state of the art.

CCS Concepts

• **Computing methodologies** → Computational photography; Image-based rendering;

1. Introduction

Radiance fields [MST*21] remain the most prevalent representation for novel-view synthesis applications due to their outstanding ability

to produce high-quality renderings of scenes from multi-view photos. 3D Gaussian representations [KKLD23] have further advanced the field and now routinely deliver real-time rendering performance.

Most works in this space are concerned with accurate, all-in-focus reconstructions of well-lit scenes. However, for radiance fields to become versatile tools for storytelling and artistic exploration, i.e., to become *cinematic*, they must simultaneously handle high-contrast illumination and depth of field. To achieve this, we propose a method for reconstructing high-dynamic-range (HDR) radiance fields with real-time re-exposure and re-focusing capabilities based on multi-view low-dynamic-range (LDR) images with varying exposures, apertures, and focus distances.

Considering the significance of the task, naturally, individual aspects of our problem setting have already been investigated in the literature. On the one hand, the limited dynamic range of typical cameras often leads to overexposed and underexposed regions, causing saturation (clamping to white) and noise, respectively [RHD*10]. These limitations directly carry over to radiance fields based on such photos and significantly restrict post-editing capabilities. Addressing this issue, several works have considered reconstructing HDR radiance fields from multi-view exposure brackets [HZF*22, JSYJYO22, RFS22] or RAW data [MHMB*22]. On the other hand, defocus blur in input images has been regarded as a degradation to overcome for all-in-focus reconstruction [WLP*22, MLL*22, LLSL23, LLS*24], with synthetic depth of field as an occasional afterthought. In contrast to all these works, we embrace both HDR and depth of field as expressive tools and set out to develop a radiance field reconstruction approach that enables real-time rendering and high-quality post-editing of exposure, aperture, and focus distance.

The input to our method is a set of LDR images that can be captured with a consumer-grade camera. Different from typical radiance field reconstruction settings, the input images vary not only in extrinsic camera parameters, but also in exposure time, aperture size, and focus distance. This diverse input sample distribution allows our method to build an all-in-focus HDR 3D Gaussian representation that is optimized for a subsequent application of lens blur. Depth of field is modeled as a highly efficient convolution of projected scene Gaussians with a Gaussian aperture function [KZB03], based on a thin-lens model. The convolved Gaussians are then rasterized and projected to LDR using a dedicated tonemapping module. Once trained, our model enables novel-view synthesis with arbitrary combinations of camera and lens parameters, facilitating flexible post-editing in terms of exposure and depth of field in real time. We show that our approach can synthesize high-quality and editable imagery, including appealing bokeh effects that require the joint treatment of HDR and depth of field. We also demonstrate that our system is on par with and often outperforms the state of the art in the individual disciplines of HDR and all-in-focus reconstruction. This evaluation is based on multiple new datasets we introduce, including both synthetic and real data, all available on our webpage <https://cinems.mpi-inf.mpg.de/>.

Our contributions can be summarized as follows:

- A unified framework for the reconstruction of HDR radiance fields with depth of field.
- Powerful post-editing capabilities such as exposure selection and arbitrary re-focusing in real time.
- State-of-the-art quality in HDR and all-in-focus reconstruction as well as new challenging benchmark datasets.

2. Related Work

In this section, we discuss previous work on depth-of-field (DoF) rendering and HDR image reconstruction. We also briefly summarize recent techniques for novel view synthesis, such as Neural Radiance Field (NeRF) and 3D Gaussian Splatting (3DGS), with special emphasis on those that support DoF and HDR effects.

2.1. Depth-of-Field Rendering

Depth-of-field (DoF) is the range within a scene where objects appear in focus, with objects outside this range appearing blurred. This effect plays an important role in photography and cinematography, helping to guide viewer attention and enhance depth perception [DG04]. Computer graphics has long sought to emulate depth-of-field effects of real-world camera systems: Cook et al. [CPC84] employ ray tracing to achieve accurate depth-of-field by casting rays across the lens. Haeberli et al. [HA90] use the accumulation-buffer technique, simulating distributed ray tracing by rendering scenes multiple times from different lens positions and blending them. Levoy et al. [LH23] apply the same concepts in light field rendering. However, the quality obtained with these methods is tied to the number of samples taken, implying a high computational effort to produce artifact-free images. To trade visual accuracy for efficiency, Scofield et al. [Sco92] apply a layered technique, segmenting scenes into depth-ordered layers for simplified blurring. This method achieves depth-of-field effects while avoiding blurred edges but struggles with depth-spanning objects. Krivanek et al. [KZB03] propose a surface splatting method suitable for real-time rendering; occasional artifacts are caused by approximate surface reconstructions. Potmesil et al. [PC82] utilize forward mapping to render sprites for depth-of-field effects and blending pixels as circles in the frame buffer, with a final renormalizing pass for producing accurate alpha values. This post-processing method for simulating depth-of-field has influenced many subsequent works [WJ22, SZW*23, PBM*22, WSP*23b], which calculate the circle of confusion (CoC) size based on depth and use spatially varying convolution to perform weighted averaging.

2.2. High Dynamic Range Image Reconstruction

Typically, we distinguish three major classes of HDR image construction [RHD*10]: specialized HDR cameras, single-shot methods that perform dynamic range expansion for an input low dynamic range (LDR) image, and multi-shot methods that merge differently exposed LDR images into an HDR output. Specialized HDR cameras have seen limited adoption in practice due to their high costs and sensor technology limitations [Eil18]. Single-shot (or *inverse tone mapping*) methods [BAD17] are an attractive alternative for handling legacy LDR images. They easily synergize with machine learning methods [MBRHD18, EKM17, LLC*20, SRK20] for denoising underexposed regions and inpainting overexposed details [EKD*17, WSP*23a]. Nevertheless, the fidelity of these methods falls short of true multi-exposure methods [DM23]. Here, multiple captures are generated at differing exposures. However, merging those exposures for dynamic scenes or handheld cameras is challenging due to ghosting and motion blur [SKY*12, HGPI2]. Pioneered by Kalantari et al. [KR*17], a further boost of reconstructed HDR

quality is achieved by learning-based methods. Recent research in this area focuses on exploring different network architectures, such as CNNs [KR*17, YGS*19, YZL*20, YZC*22, XC21] and Transformers [CZZ*23, LWZL22, TWZ*23]. Further work explores the integration of generative adversarial losses [NWL*21], diffusion priors [YHS*23], as well as semi-supervised [YZC*23, PSA*21] and self-supervised methodologies [ZWL*24].

2.3. Depth-of-Field and High Dynamic Range Radiance Fields

Significant progress has been made in reconstructing radiance fields from 2D images, with NeRF [MST*21] and 3D Gaussian Splatting [KKLD23] techniques being particularly noteworthy. Neural Radiance Fields (NeRF) [MST*21] use a multi-layer perceptron (MLP) with positional encoding to map 3D locations and viewing directions to color and density values. Parameter optimization is achieved via gradient descent on the photometric error. Subsequent work on NeRF [HSM*21, VHM*22, BMT*21, BMV*22, BMV*23, MESK22] focused on further improving its quality and initially slow rendering speed of this implicit representation. In contrast, 3D Gaussian Splatting (3DGS) [KKLD23] uses explicit 3D Gaussian primitives as a flexible and expressive scene representation. The optimization starts from point clouds and optimizes per-Gaussian properties. While effective, the recency of 3DGS entails minor flaws in its method. Mip-Splatting [YCH*24] and multi-scale 3D Gaussian splatting [YLCL24] both address aliasing artifacts present in the original approach. Mip-Splatting uses 3D size constraints and 2D Mip filters to enhance image quality for zoomed-in or zoomed-out views. Multi-scale 3D Gaussian splatting maintains Gaussians at different scales, using more small Gaussians for high resolution and larger Gaussians for low resolution.

In most radiance field research, it is assumed that input images are consistently in focus and well-exposed. While many methods can tolerate small deviations from those assumptions, few recent methods more actively model variations in focal length and exposure time. DoF-NeRF [WLP*22] extends NeRF to handle shallow depth-of-field inputs and simulate DoF effects through a differentiable circle of confusion representation. Deblur-NeRF [MLL*22] improves NeRF by addressing image blurriness from defocus or motion with a Deformable Sparse Kernel (DSK) module, enabling sharp 3D reconstructions. DP-NeRF [LLSL23] tackles geometric and appearance consistency issues in blurred images using physical priors and adaptive weighting, significantly enhancing 3D reconstruction quality in the presence of motion and defocus blur. Concurrent work [LLS*24, DPRBK24] proposes to compensate for defocus blur in captures with Gaussian splats by inferring alternative covariances for out-of-focus Gaussians. HDR-NeRF [HZF*22] recovers HDR radiance fields from LDR views with different exposures, enabling the generation of novel HDR and LDR views by modeling the physical imaging process integrated with NeRF. In freshly appearing concurrent work [CXL*24], a similar approach has been proposed for 3DGS that greatly reduces training and inference time. RawNeRF [SGHS98, MHMB*22] utilizes a modified NeRF technique that trains directly on linear raw images to preserve the scene's full dynamic range, enabling high-quality HDR novel view synthesis. It should be noted that RawNeRF uses raw format inputs without defocus blur and requires the multiplane image-based

blurring algorithm [KS07] for refocusing. In contrast, our approach uses defocused 8-bit LDR inputs, jointly models depth of field and tone mapping processes, and inherently supports all post-editing tasks. The methods discussed above can handle simple DoF and HDR factors independently. However, they cannot account for both simultaneously; this causes issues in casual captures, e.g., scenes that contain details in the near foreground and far background, as well as strong luminance variation between dark shadows and bright lights.

As shown by TAF [WSP*23b], a stack of images with different exposure/aperture/focus can be employed for all-in-focus HDR image reconstruction of particularly demanding scenes for a specific camera pose. Our goal is to generalize this approach to multiple camera poses within a novel-view synthesis framework based on 3DGS for efficient rendering. Capturing a complete image stack for each input camera pose would be impractical; we aim to sample the space of camera parameter setups between neighboring camera poses so that, as in traditional NeRF and 3DGS methods, a single image per pose is effectively captured. This way, we enable high-quality reconstruction of scenes with high variation of depth and contrast ranges, and their rendering with arbitrary camera parameters for any view. Our approach directly benefits from the more realistic depth-of-field simulations when performed using HDR images as the input [ZMN*19].

3. Method

Our approach takes as input a set of LDR multi-view images $\{I_k\}_{k=1}^N$ of a scene, captured with randomly selected exposure times t , aperture sizes a , and focus distances d_f (Fig. 1, left). Since exposure and lens settings are typically stored in EXIF headers, we assume this information is available for each image. Using this input, our method reconstructs an HDR radiance field that allows for post-editing of exposure and depth of field.

Our approach is based on the 3D Gaussian Splatting representation [KKLD23], which we briefly review in Sec. 3.1. The first step in our pipeline is calibrating the input cameras, detailed in Sec. 3.2. We then introduce our model, which features an HDR 3D Gaussian scene representation with an integrated defocus blur simulator (Sec. 3.3). A tonemapping module allows us to train the representation on LDR input only (Sec. 3.4).

3.1. Background

Radiance fields based on 3D Gaussians [KKLD23] represent a scene using a Gaussian mixture model. Each primitive has a mean $\mu \in \mathbb{R}^3$, a symmetric covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$, an opacity $o \in \mathbb{R}$, as well as a set of spherical harmonics (SH) coefficients encoding view-dependent RGB color $\mathbf{c} \in \mathbb{R}^3$. To render an image using this representation [ZPVBG01], the primitives are projected to 2D screen space using a pinhole camera model to obtain a projected mean $\hat{\mu} \in \mathbb{R}^2$ and projected covariance $\hat{\Sigma} \in \mathbb{R}^{2 \times 2}$. The final pixel color is obtained using front-to-back compositing:

$$C(\mathbf{x}) = \sum_{i \in \mathcal{N}(\mathbf{x})} \mathbf{c}_i \alpha_i(\mathbf{x}) \prod_{j=1}^{i-1} (1 - \alpha_j(\mathbf{x})), \quad (1)$$

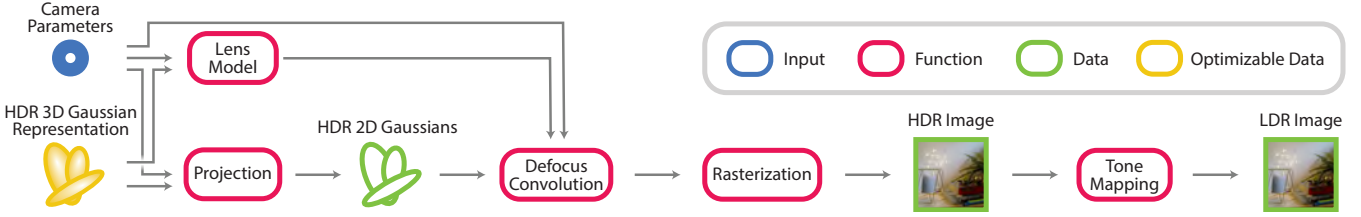


Figure 2: Our pipeline for rendering HDR radiance fields with controllable depth of field.

where \mathcal{N} is the sequence of depth-ordered primitives overlapping the pixel at location $\mathbf{x} \in \mathbb{R}^2$, and

$$\alpha_i(\mathbf{x}) = o_i \cdot \exp\left(-\frac{1}{2}(\mathbf{x} - \hat{\mu}_i)^T \hat{\Sigma}_i^{-1}(\mathbf{x} - \hat{\mu}_i)\right) \quad (2)$$

is the opacity of the primitives. Since the rendering function in Eq. 1 is differentiable, the model can be trained from posed input images via gradient-based optimization.

3.2. Camera Calibration

We require posed input images to build our model. The standard approach for this is Structure from Motion (SfM) [SF16], which we run on our input data. Unfortunately, such photometric calibration leaves the scene's global scale ambiguous. As various lens parameters are provided in absolute millimeters, we need to scale the scene accordingly. To achieve this, we leverage the defocused images and their respective lens configurations. Specifically, we run an image blur detector [AGK17] on all input images to obtain per-pixel sharpness estimates. By thresholding these sharpness maps, we identify the image space locations of in-focus scene elements. The known focus distance d_f per input image provides ground-truth depth values for these pixels. Next, we train a vanilla 3D Gaussian Splatting model on the input images after rescaling pixel intensities to a common exposure. Using this model, we render a depth map for each input view. As expected, the renderings from this model are of low quality given our inputs, but we find the overall scene configuration precise enough for our calibration. Using the depth values of the in-focus pixels, we perform a least-squares fit of the scene scale to match the rendered depth values with the ground truth.

After calibration, we have the full set of camera and lens parameters per input view at our disposal. Please note that we follow the methodology established in 3DGS [KKLD23]. The initialization process differs for synthetic and real datasets: for real scenes, we use SfM for camera and scale calibration, and initialize the scene Gaussians using SfM points. For synthetic scenes with predefined camera parameters, SfM is unnecessary, eliminating scale ambiguity and calibration. In this case, scene Gaussians are initialized with random noise (see Sec. 4.1.2).

3.3. HDR Radiance Field with Depth of Field

Here we describe our pipeline, which consists of an HDR 3D Gaussian scene representation, a defocus convolution module, as well as

a tonemapper that projects HDR to LDR content. An overview of our model is shown in Fig. 2.

Our HDR Gaussians differ from conventional LDR Gaussians in one crucial property: the per-primitive SH coefficients encoding view-dependent color \mathbf{c} exhibit a significantly larger dynamic range. We denote this corresponding HDR color as \mathbf{c}^{HDR} . Notably, no architectural changes compared to an LDR representation are necessary at this stage; the high dynamic range emerges from our carefully designed pipeline, as detailed in the remainder of this section.

The next stage of our model is concerned with depth of field (DoF) synthesis, which occurs at the level of 2D Gaussians projected into screen space. Considering a thin-lens model [PC81], a projected point with depth d is spread over a disk – typically referred to as the circle of confusion – with radius

$$r(d) = \frac{|d - d_f| \cdot f^2}{2 \cdot a \cdot d \cdot (d_f - f)}, \quad (3)$$

where f is the focal length, a is the F-number, representing the aperture size and d_f is the focus distance. We seek to enlarge each projected scene Gaussian to reflect the circle of confusion [KZB03], where the mean μ determines depth d . We approximate the circle of confusion using an aperture Gaussian with covariance $\hat{\Sigma}^a = \text{diag}(3r(d), 3r(d)) \in \mathbb{R}^{2 \times 2}$ to allow for efficient convolutions with the scene Gaussians. This convolution can be carried out in closed form [Bro03, CHK*22], resulting in an update of Eq. 2:

$$\alpha_i^{\text{DoF}}(\mathbf{x}) = o_i \cdot \beta_i \cdot \exp\left(-\frac{1}{2}(\mathbf{x} - \hat{\mu}_i)^T (\hat{\Sigma}_i + \hat{\Sigma}_i^a)^{-1}(\mathbf{x} - \hat{\mu}_i)\right), \quad (4)$$

with the energy-preserving normalization factor

$$\beta_i = \sqrt{\frac{\det(\hat{\Sigma}_i)}{\det(\hat{\Sigma}_i + \hat{\Sigma}_i^a)}}. \quad (5)$$

Importantly, each projected scene Gaussian is convolved with an aperture Gaussian reflecting the depth of the scene Gaussian and the lens configuration, producing spatially-varying defocus blur. An HDR image with defocus blur can now be rasterized using the quantities developed above:

$$C_{\text{HDR+DoF}}(\mathbf{x}) = \sum_{i \in \mathcal{N}(\mathbf{x})} \mathbf{c}_i^{\text{HDR}} \alpha_i^{\text{DoF}}(\mathbf{x}) \prod_{j=1}^{i-1} (1 - \alpha_j^{\text{DoF}}(\mathbf{x})). \quad (6)$$

Naturally, the HDR content rendered in the previous step has an unbounded range of radiance values. Both training and final visualization require values in the range $[0, 1]$. Therefore, we project

HDR values to LDR using the tone mapping function

$$\bar{C}(\mathbf{x}) = \left(\min\left(\frac{t}{a^2} \cdot C_{\text{HDR+DoF}}(\mathbf{x}), 1\right) \right)^{1/\gamma}, \quad (7)$$

simulating the per-pixel non-linear response of a real-world camera [AT11]. Here, the final exposure is determined using exposure time t and aperture size a , before the intensities are clipped to the target domain. Finally, γ performs a gamma correction from linear radiance to pixel intensities. We use an explicit tone mapping operator as it minimizes computational overhead and allows convenient metadata incorporation.

Eq. 7 produces images that the user can flexibly edit across multiple dimensions: In addition to free viewpoint selection, the user can adjust the aperture size a , exposure time t , and focus distance d_f in real-time, providing a wide range of artistic control.

3.4. Training

We train our model using the loss function

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \lambda_{\text{exp}} \mathcal{L}_{\text{exp}} + \lambda_{\text{foc}} \mathcal{L}_{\text{foc}}. \quad (8)$$

All three terms make use of the photometric consistency measure Δ that evaluates a sum of ℓ_1 -distance and DSSIM [WBSS04] metric. Yet, each term applies Δ to different versions of the data. The first term simply measures the consistency between our rendered images and the input data for all pixels \mathbf{x} and training views k :

$$\mathcal{L}_{\text{rec}} = \mathbb{E}_{\mathbf{x},k} [\Delta(\bar{C}_k(\mathbf{x}), I_k(\mathbf{x}))]. \quad (9)$$

To handle the vastly different exposures in our inputs, the second loss term normalizes the images to a medium exposure by scaling pixel intensities before computing discrepancies [LLC*20, WCS*22] (Fig. 3, left and center):

$$\mathcal{L}_{\text{exp}} = \mathbb{E}_{\mathbf{x},k} \left[\Delta \left(\frac{\bar{C}_k(\mathbf{x})}{2 \cdot \mathbb{E}_{\mathbf{x}}[\bar{C}_k(\mathbf{x})]}, \frac{I_k(\mathbf{x})}{2 \cdot \mathbb{E}_{\mathbf{x}}[I_k(\mathbf{x})]} \right) \right]. \quad (10)$$

Finally, due to the prominent occurrence of defocus blur in our inputs, we found it beneficial to include a loss term that explicitly addresses sharpness. Specifically, we use an off-the-shelf differentiable image blur detector g [AGK17] that outputs a scalar defocus map (Fig. 3, right) and compute

$$\mathcal{L}_{\text{foc}} = \mathbb{E}_{\mathbf{x},k} [\Delta(g(\bar{C}_k(\mathbf{x})), g(I_k(\mathbf{x})))] . \quad (11)$$

We set $\lambda_{\text{exp}} = 0.25$ and $\lambda_{\text{foc}} = 1$ in all our experiments and train our models using the Adam [KB15] optimizer with default parameters, also including all densification strategies from 3DGS [KKLD23].

4. Results and Evaluation

In this section, we demonstrate the effectiveness of our method through a comprehensive set of experiments. We first describe our experimental settings in Sec. 4.1, introducing our dataset, which consists of both synthetic and real captured images used for evaluation, followed by the implementation details used for all results. Then, in Sec. 4.2, we evaluate separately the capability of our model to reconstruct all-in-focus and HDR images, comparing our results to those of specialized state-of-the-art solutions. After this, we present




Figure 3: Our training data includes underexposed and overexposed images (left). Our loss incorporates both the original data (Eq. 9) and a normalized version of it (Eq. 10). To manage defocus blur, a third loss component (Eq. 11) incorporates a defocus map.

in Sec. 4.3 results for our complete pipeline, showcasing HDR radiance field reconstructions as well as post-editing applications. Lastly, in Sec. 4.4 we provide ablation studies of the components of our approach. Please refer to the supplementary video for additional results.

4.1. Experimental Settings

4.1.1. Dataset

Our dataset includes four synthetic rendered scenes and four real captured scenes. The synthetic scenes provide control and ground truth for evaluation, while the real captured scenes demonstrate the practical performance of our approach. A sample of all scenes can be found in the supplementary material.

Synthetic dataset Our synthetic dataset is composed of four diverse indoor scenes rendered with Cycles (path tracer) in Blender 4.1. The scenes encompass challenging scenarios with high dynamic range of illumination (bright light sources and dark regions illuminated by bounced light) and spacious layouts allowing for multiple focal planes (examples shown in Figures 5 and 7). All scenes have been rendered with high number of samples (ranging from 4096 to 16384) to prevent undesired distortions in the defocused regions caused by denoising. For each scene we generated 960 HDR images (8×8 views sampled at a regular grid wrapped to the sphere surface, 5 uniformly sampled focus distances and 3 different apertures). For each view we generated all-in-focus ground truth image. This results in 1024 images per scene. For evaluation purposes we generated another set of 7×7 views located between the training views. All images are accompanied by a file containing parameters used for rendering (camera transformation matrix, focal length, focus distance, aperture, etc.). In addition to the HDR content we have modified the scenes for the AiF reconstruction comparison, producing pleasant LDR images with reduced dynamic range without clipped regions (examples shown in Figure 4). The camera locations, apertures, and focus distances are sampled in a same manner as in the case of the HDR subset.

Captured dataset Our captured dataset is composed of four real scenes. Some examples are shown in Figs. 1, 6, and 7. These scenes were captured using Canon EOS RP and Canon EOS Rebel T6i cameras, outfitted with RF 85mm F2.0, and EF-S 18-55mm F3.5-5.6 lenses. Each scene was captured by moving the camera to follow an approximately 9×9 grid pattern, resulting in 81 images per scene.

At each camera position in this grid, we varied the exposure time, aperture, and focus distance to ensure a diverse set of configurations.

4.1.2. Implementation

We use the 3DGS framework [KKLD23] as our baseline, integrating a thin-lens model and a tone mapping module as described in Sec. 3.3. The implementation is done in PyTorch and CUDA. All source code is available on our project webpage.

Additionally, we incorporate some practical design choices to tackle our task more effectively. Given that our synthetic dataset starts from random noise without an SfM point cloud initialization, we employ a coarse-to-fine optimization strategy. During the initial 7,000 iterations, we exclude tone mapping and depth of field simulation, training on an exposure-aligned dataset to generate an initial point cloud. Following this, we refine the optimization using the coarse result as a starting point. Our first stage training adapts well to a broad range of dynamic variations. However, due to sparse viewpoints, we observed instability when the scene's dynamic range exceeds 18 stops. The iteration process consists of 40,000 steps, taking approximately 21 minutes on an RTX A40 GPU, using a training dataset with a resolution of 1200×675 .

4.2. Comparisons

In this section we evaluate two important features of our method: all-in-focus reconstruction and HDR reconstruction. We compare our method against specialized state-of-the-art approaches designed to address these specific problems. To facilitate this comparison, we leverage our synthetic dataset, which provides access to ground truth images for numerical evaluations. While these specialized methods excel in their respective domains, it is important to note that they are limited to solving only one of the problems—those optimized for all-in-focus reconstruction cannot handle effectively HDR, and vice versa. Our approach, however, integrates both capabilities and demonstrates comparable, and in some cases superior performance to these specialized solutions.

4.2.1. All-in-Focus Reconstruction

For the all-in-focus reconstruction task, we compare our results against the concurrent work Deblur-Splatting [LLS*24]. Additionally, we include in the supplementary material comparisons for Deblur NeRF [MLL*22], where we observed noticeable deformations along the boundaries, causing straight lines to appear curved. We use the four scenes from our synthetic dataset for this evaluation. In particular, for each of the scenes, we use 64 (8×8) viewpoints for training, each exhibiting varying degrees of defocus blur, and 49 (7×7) novel viewpoints not seen during training for the evaluation. Table 1 presents the numerical results for PSNR, SSIM [WBSS04], and LPIPS [ZIE*18], commonly used for evaluating image quality. Additionally, Fig. 4 shows illustrative comparisons to the ground truth. Although our method is not specifically designed for deblurring, it achieves results comparable to state-of-the-art method. Regarding rendering time, both methods are similarly efficient, achieving approximately 110 frames per second (FPS) on an RTX A40 GPU. We primarily evaluate the synthetic dataset due to the challenge of obtaining ground truth for real datasets. To validate our method, we

run Deblur-Splatting on our real dataset using non-reference metrics. Since Deblur-Splatting cannot generate HDR, we tone-map our HDR images and evaluate in the LDR domain using NIQE [MSB12] and BRISQUE [MMB12]. The results in Table 2 demonstrate our method's superior performance.

Table 1: Comparison of all-in-focus reconstruction performance metrics for our method and Deblur-Splatting using our synthetic dataset.

Scene	Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Car	Deblur-Splatting	27.23 \pm 3.21	0.86 \pm 0.07	0.18 \pm 0.04
	Ours	28.05 \pm 0.26	0.89 \pm 0.01	0.11 \pm 0.01
Attic	Deblur-Splatting	31.86 \pm 0.23	0.91 \pm 0.01	0.13 \pm 0.01
	Ours	32.90 \pm 0.18	0.93 \pm 0.00	0.09 \pm 0.00
Café	Deblur-Splatting	30.24 \pm 0.42	0.87 \pm 0.01	0.12 \pm 0.01
	Ours	30.50 \pm 0.36	0.85 \pm 0.01	0.14 \pm 0.00
Office	Deblur-Splatting	34.59 \pm 0.75	0.95 \pm 0.01	0.08 \pm 0.01
	Ours	31.70 \pm 0.80	0.92 \pm 0.01	0.05 \pm 0.01

4.2.2. HDR Reconstruction

For HDR reconstruction, we compare our results against HDR-NeRF [HZF*22] using various metrics specifically targeted towards evaluating HDR quality, including HDR-VDP-3 [MHH23], PU-PSNR, and PU-SSIM [MA21], as summarized in Table 3. Visual comparisons with the ground truth are depicted in Fig. 5. Again, we use for evaluation the four scenes of our synthetic dataset, with 64 (8×8) views for training and 49 (7×7) novel ones for testing. To simulate LDR training views with varying exposures, starting from the HDR renders from our synthetic dataset, we sample a single exposure for each view. For this, we first determine the start and end boundaries that encompass the entire dynamic range of the scene [ANSAM21]. Then, we divide this range into eight intervals, within which we sample exposures as described in previous work [HZF*22]. Our method consistently outperforms HDR-NeRF in almost all metrics, demonstrating superior reconstruction capabilities. Notably, we observed that HDR-NeRF tends to overfit to the LDR images, resulting in lower quality HDR image generation. Moreover, our approach achieves real-time rendering capabilities, significantly outperforming HDR-NeRF. We deliver approximately 110 FPS at a resolution of 1200×675 , in contrast to HDR-NeRF's rate of approximately 0.021 FPS for rendering a single image, both using a GPU RTX A40. We also run HDR-NeRF on our real dataset using non-reference metrics PU-PIQE [HME*22] to assess results from novel views (Table 4). These results highlight that our method outperforms task-specific approaches.

Table 2: Comparison of all-in-focus reconstruction performance metrics for our method and Deblur-Splatting using our real dataset.

Method	NIQE \downarrow	BRISQUE \downarrow
Deblur-Splatting	6.65 \pm 1.87	45.77 \pm 8.50
Ours	2.66 \pm 0.22	32.17 \pm 4.75




Figure 4: Comparison of all-in-focus reconstruction. Our method provides better or comparable results to Deblur-Splatting method.




Figure 5: Comparison of HDR reconstruction. Left-most column presents input LDR images for a given view. The following columns depict results acquired by HDR-NeRF, Ground Truth, and results of our method, respectively. Please note, that our method properly reconstructs the shape of the flame and reproduces successfully high-frequency content, while these details are missing in the results generated by HDR-NeRF. To depict the accuracy of both methods we provide the HDR-VDP-3 error map in the top right corner of the images. To better show the differences in the highlights, we render these regions at -3 stops, as shown in the yellow box.

4.3. HDR Radiance Field Reconstruction and Post-Editing

In this section, we showcase the full capabilities of our model, specifically its ability to reconstruct HDR radiance fields, which enable the generation of all-in-focus HDR images. Fig. 6 presents results from various viewpoints, illustrating how our approach effectively reconstructs detailed scenes with consistent sharpness and dynamic range from multi-view LDR defocused inputs. Furthermore, we demonstrate the post-editing capabilities enabled by our method. Our proposed approach reconstructs HDR radiance fields with real-time re-exposure and re-focusing capabilities. Users have the flexibility to adjust aperture size, exposure time, and focus distance in real-time, as demonstrated in Fig. 7.

4.4. Ablation Study

We evaluate the impact of the defocus and exposure loss terms by conducting ablation studies on both our real and synthetic datasets. Additionally, we include in the supplementary material an evaluation of our coarse-to-fine training strategy for the synthetic dataset.

We conduct this evaluation differently for our synthetic and real datasets. For the synthetic dataset, where ground-truth images are available, we evaluate our all-in-focus HDR reconstructions against ground-truth views not included in the training set. Specifically, we utilize 64 (8×8) views for training and 49 (7×7) novel views for testing. This comparison is performed in HDR space, therefore we use PU-PSNR, PU-SSIM and HDR-VDP-3 as metrics. We show these results in Table 5. On the other hand, for the captured dataset, where ground truth is not available, we adopt a different approach. We use 74 captured views for training and leave out 7 for testing. Then, we synthesize the testing views with our model

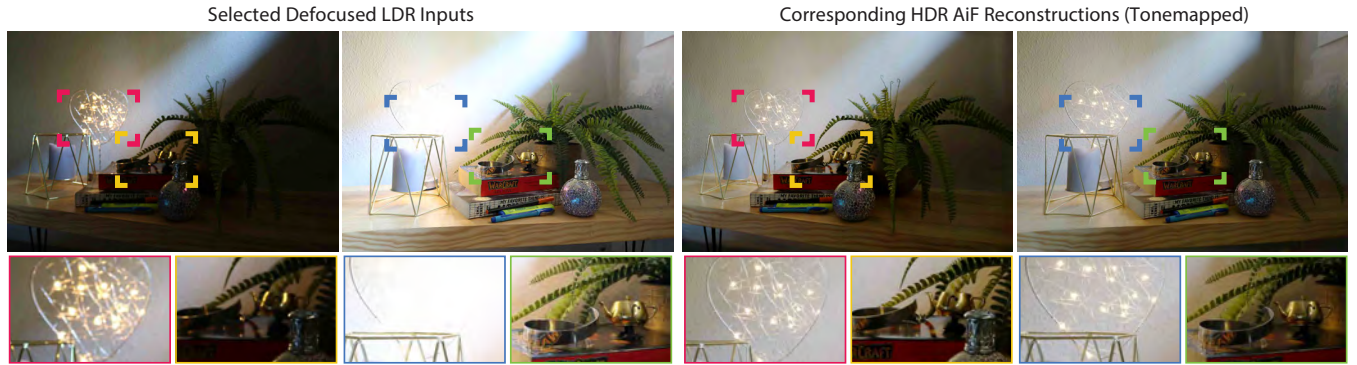


Figure 6: All-in-focus HDR reconstruction. Our method can successfully reconstruct sharp HDR images from defocused images captured with different exposure times.

Table 3: Comparison of HDR reconstruction performance metrics for our method and HDR-NeRF using our synthetic dataset.

Scene	Method	HDR-VDP-3 \uparrow	PU-PSNR \uparrow	PU-SSIM \uparrow
Car	HDR-NeRF	9.70 \pm 0.15	29.78 \pm 1.72	0.91 \pm 0.02
	Ours	9.75 \pm 0.09	29.67 \pm 4.63	0.95 \pm 0.01
Attic	HDR-NeRF	9.92 \pm 0.03	39.04 \pm 0.71	0.96 \pm 0.04
	Ours	9.98 \pm 0.01	38.97 \pm 1.01	0.98 \pm 0.00
Café	HDR-NeRF	9.16 \pm 0.22	26.97 \pm 2.56	0.76 \pm 0.11
	Ours	9.93 \pm 0.01	30.02 \pm 0.51	0.91 \pm 0.00
Office	HDR-NeRF	9.36 \pm 0.18	25.13 \pm 1.97	0.90 \pm 0.01
	Ours	9.50 \pm 0.09	30.11 \pm 0.22	0.71 \pm 0.05

Table 4: Comparison of HDR reconstruction performance metric for our method and HDR-NeRF using our real dataset.

Method	PU-PIQE \downarrow
HDR-NeRF	55.47 \pm 2.22
Ours	42.22 \pm 1.22

using the original aperture, focus distance, and exposure settings corresponding to each captured view, and compare these synthesized views against the actual captures. This comparison is conducted in LDR space, thus we employ traditional image quality metrics, in particular, PSNR, SSIM, and LPIPS. Quantitative results are presented in Table 6, while qualitative results are illustrated in Fig. 8 for the defocus loss term and in Fig. 9 for the exposure loss term.

Table 5: Evaluation of the impact of the defocus and exposure loss terms in the synthetic dataset.

Method	PU-PSNR \uparrow	PU-SSIM \uparrow	HDR-VDP \uparrow
w/o \mathcal{L}_{exp}	33.95 \pm 1.59	0.94 \pm 0.00	9.80 \pm 0.06
w/o \mathcal{L}_{foc}	33.63 \pm 1.54	0.94 \pm 0.00	9.78 \pm 0.05
Ours	34.04 \pm 1.67	0.95 \pm 0.00	9.83 \pm 0.05

Table 6: Evaluation of the impact of the defocus and exposure loss terms in the real dataset.

Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o \mathcal{L}_{exp}	23.95 \pm 3.37	0.82 \pm 0.10	0.11 \pm 0.04
w/o \mathcal{L}_{foc}	24.35 \pm 1.06	0.85 \pm 0.06	0.11 \pm 0.03
Ours	25.69 \pm 2.71	0.86 \pm 0.07	0.10 \pm 0.02

4.5. Discussion of Depth-of-Field Generation

Our depth-of-field module can achieve real-time performance at high quality. Alternatives include accumulation buffering [HA90], which delivers high-quality results but is slow due to the need for multiple samples. Another method, used in RawNeRF [MHMB*22], decomposes the image into multiplane images at different depths [SGHS98, KS07], applies image space blurring, and then composites them. While this improves speed, it struggles with large depth complexity. In contrast, our approach does not have this limitation.

To shed some light on the respective trade-offs, we compare our method against accumulation buffering for depth of field simulation. Specifically, we replace our DoF module with accumulation buffering, testing with 20 (AB_{low}) and 120 samples (AB_{high}). Unlike accumulation buffering, which depends on Monte Carlo sampling, our method requires only a single rendering. As shown in Table 7, our method delivers depth-of-field quality comparable to accumulation buffering with a high sample count while significantly improving speed, enabling real-time rendering. Visual comparisons in Fig. 10 reveal that lower sampling in accumulation buffering results in noticeable artifacts.

Table 7: Comparisons of our depth-of-field rendering module against solutions involving an accumulation buffer.

Method	PSNR \uparrow	SSIM \uparrow	FPS \uparrow
AB_{high}	27.59 \pm 3.39	0.94 \pm 0.02	1
AB_{low}	27.29 \pm 3.34	0.93 \pm 0.02	6
Ours	27.41 \pm 3.30	0.94 \pm 0.02	90




Figure 7: Post-editing results. The leftmost column shows the selected input views. The 3×3 grid to the right displays novel views, with each row illustrating the manipulation capabilities of one parameter: focus distance, aperture size, and exposure.




Figure 8: The defocus loss provides higher quality reconstruction for high-frequency patterns.



Figure 9: The exposure loss helps to resolve artifacts in the under-exposed regions, as the one depicted in the green frame.

5. Limitations and Future Work

Our approach does not come without limitations, opening up fruitful avenues for future work. Our defocus model is a two-fold approximation of the actual physical processes in a real camera lens. First, we use a thin-lens model [PC81], which, despite being commonly used, is a simplified representation of defocus. More sophisticated models have been developed in the forward-rendering literature [KMH95] and related work on defocus modeling in the context of reconstruction from image stacks [WSP*23b] has shown that advanced effects, such as lens breathing, can significantly improve accuracy. Second, we approximate the thin-lens circle of confusion with a Gaussian function to achieve real-time rendering performance. While this is a commonly used design decision [Fav10, SZW*23], it results in a slight softening of bokeh due to the deviation from a disk kernel. We believe that moving towards more realistic lens models has the potential to further increase result quality and cinematic appeal.




Figure 10: Depth-of-field rendering comparison. Our real-time method achieves high-quality results comparable to those using an accumulation buffer with many samples.



Figure 11: Limitation. All-in-focus HDR reconstruction of a scene view with extremely high dynamic range, featuring tone-mapped low (left) and high (right) exposures. In regions with pronounced darkness and notable noise in the input images, such as the rear wall of the room, artifacts may arise due to information loss.

In scenes characterized by extremely high dynamic range, where very dark regions may contain significant noise across most input images, artifacts may occur in the reconstruction of these dark areas. These artifacts are pronounced when selecting very high exposures (Fig. 11). This is expected, since there is insufficient useful information, with noise dominating instead in these areas. Nevertheless, these artifacts are noticeable only under extreme exposure settings in the reconstruction, while our method remains capable of delivering pleasing results for the rest of the scene. As discussed above, reduced exposure, whether through shorter exposure times or a smaller aperture, results in increased noise in real camera images. However, this paper primarily focuses on HDR and depth of field characteristics. Addressing noise in dark regions is identified as a potential area for future research. Therefore, for our synthetic dataset, we adhere to the HDR-Nerf [HZF*22] procedure to render noise-free images.

6. Conclusion

We have presented an approach for reconstructing HDR radiance fields with depth of field. Unlike previous works that treat only exposure variations or defocus blur as a degradation to overcome, we recognize the significance of exposure selection and depth of field for crafting cinematic imagery. Our pipeline mirrors the physical processes of a camera, enabling us to train a radiance field model using LDR images with various exposure and lens configurations. We found that our method not only allows powerful and appealing post-editing in real time, but also achieves state-of-the-art quality for the isolated tasks of HDR and all-in-focus reconstruction. We hope that our approach moves radiance fields a step closer to becoming expressive tools for artists and visual media producers.

Acknowledgements

This work has been partially supported by grant PID2022-141539NB-I00, funded by MICIU/AEI/10.13039/501100011033 and by ERDF, EU. Open Access funding enabled and organized by Projekt DEAL.

References

- [AGK17] ALIREZA GOLESTANEH S., KARAM L. J.: Spatially-varying blur detection based on multiscale fused and sorted transform coefficients of gradient magnitudes. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017), pp. 5800–5809. 4, 5

- [ANSAM21] ANDERSSON P., NILSSON J., SHIRLEY P., AKENINE-MÖLLER T.: Visualizing errors in rendered high dynamic range images. 6
- [AT11] ALLEN E., TRIANTAPHILLIDOU S.: *The manual of photography*. Taylor & Francis, 2011. 5
- [BADC17] BANTERLE F., ARTUSI A., DEBATTISTA K., CHALMERS A.: *Advanced high dynamic range imaging*. AK Peters/CRC Press, 2017. 2
- [BMT*21] BARRON J. T., MILDENHALL B., TANCİK M., HEDMAN P., MARTIN-BRUALLA R., SRINIVASAN P. P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (2021)*, pp. 5855–5864. 3
- [BMV*22] BARRON J. T., MILDENHALL B., VERBIN D., SRINIVASAN P. P., HEDMAN P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022)*, pp. 5470–5479. 3
- [BMV*23] BARRON J. T., MILDENHALL B., VERBIN D., SRINIVASAN P. P., HEDMAN P.: Zip-nerf: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (2023)*, pp. 19697–19705. 3
- [Bro03] BROMILEY P.: Products and convolutions of gaussian probability density functions. *Tina-Vision Memo* 3, 4 (2003), 1. 4
- [CHK*22] CELAREK A., HERMOSILLA P., KERBL B., ROPINSKI T., WIMMER M.: Gaussian mixture convolution networks. In *International Conference on Learning Representations (2022)*. 4
- [CPC84] COOK R. L., PORTER T., CARPENTER L.: Distributed ray tracing. In *Proceedings of the 11th annual conference on Computer graphics and interactive techniques (1984)*, pp. 137–145. 2
- [CXL*24] CAI Y., XIAO Z., LIANG Y., QIN M., ZHANG Y., YANG X., LIU Y., YUILLE A.: HDR-GS: Efficient high dynamic range novel view synthesis at 1000x speed via gaussian splatting, 2024. [arXiv: 2405.15125](https://arxiv.org/abs/2405.15125). 3
- [CZZ*23] CHEN R., ZHENG B., ZHANG H., CHEN Q., YAN C., SLABAUGH G., YUAN S.: Improving dynamic hdr imaging with fusion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence (2023)*, vol. 37, pp. 340–349. 3
- [DG04] DEMERS J., GEMS G.: Chapter 23. depth of field: A survey of techniques. *CPU Gems, Developer Zone, copyright (2004)*. 2
- [DM23] DEBEVEC P. E., MALIK J.: Recovering high dynamic range radiance maps from photographs. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 2023, pp. 643–652. 2
- [DPRBK24] DARMON F., PORZI L., ROTA-BULÒ S., KONTSCIEDER P.: Robust gaussian splatting, 2024. [arXiv:2404.04211](https://arxiv.org/abs/2404.04211). 3
- [Eil18] EILERTSEN G.: *The high dynamic range imaging pipeline*, vol. 1939. Linköping University Electronic Press, 2018. 2
- [EKD*17] EILERTSEN G., KRONANDER J., DENES G., MANTIUK R. K., UNGER J.: Hdr image reconstruction from a single exposure using deep cnns. *ACM transactions on graphics (TOG)* 36, 6 (2017), 1–15. 2
- [EKM17] ENDO Y., KANAMORI Y., MITANI J.: Deep reverse tone mapping. *ACM Trans. Graph.* 36, 6 (2017), 177–1. 2
- [Fav10] FAVARO P.: Recovering thin structures via nonlocal-means regularization with application to depth from defocus. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2010)*, IEEE, pp. 1133–1140. 10
- [HA90] HAEBERLI P., AKELEY K.: The accumulation buffer: Hardware support for high-quality rendering. *ACM SIGGRAPH computer graphics* 24, 4 (1990), 309–318. 2, 8
- [HGP12] HU J., GALLO O., PULLI K.: Exposure stacks of live scenes with hand-held cameras. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part I 12 (2012)*, Springer, pp. 499–512. 2
- [HME*22] HANJI P., MANTIUK R., EILERTSEN G., HAJISHARIF S., UNGER J.: Comparison of single image hdr reconstruction methods—the caveats of quality assessment. In *ACM SIGGRAPH 2022 conference proceedings (2022)*, pp. 1–8. 6
- [HSM*21] HEDMAN P., SRINIVASAN P. P., MILDENHALL B., BARRON J. T., DEBEVEC P.: Baking neural radiance fields for real-time view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (2021)*, pp. 5875–5884. 3
- [HZF*22] HUANG X., ZHANG Q., FENG Y., LI H., WANG X., WANG Q.: HDR-NeRF: High dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022)*, pp. 18398–18408. 2, 3, 6, 10
- [JSYJYBO22] JUN-SEONG K., YU-JI K., YE-BIN M., OH T.-H.: HDR-pixeloxels: Self-calibrating high dynamic range radiance fields. In *European Conference on Computer Vision (2022)*, Springer, pp. 384–401. 2
- [KB15] KINGMA D. P., BA J.: Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR) (2015)*. 5
- [KKLD23] KERBL B., KOPANAS G., LEIMKÜHLER T., DRETTAKIS G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* 42, 4 (2023), 1–14. 1, 3, 4, 5, 6
- [KMH95] KOLB C., MITCHELL D., HANRAHAN P.: A realistic camera model for computer graphics. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques (1995)*, pp. 317–324. 10
- [KR*17] KALANTARI N. K., RAMAMOORTHI R., ET AL.: Deep high dynamic range imaging of dynamic scenes. *ACM Trans. Graph.* 36, 4 (2017), 144–1. 2, 3
- [KS07] KRAUS M., STRENGERT M.: Depth-of-field rendering by pyramidal image processing. In *Computer graphics forum (2007)*, vol. 26, Wiley Online Library, pp. 645–654. 3, 8
- [KZB03] KRIVÁNEK J., ZARA J., BOUATOUCH K.: Fast depth of field rendering with surface splatting. In *Proceedings Computer Graphics International 2003 (2003)*, IEEE, pp. 196–201. 2, 4
- [LH23] LEVOY M., HANRAHAN P.: Light field rendering. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 2023, pp. 441–452. 2
- [LLC*20] LIU Y.-L., LAI W.-S., CHEN Y.-S., KAO Y.-L., YANG M.-H., CHUANG Y.-Y., HUANG J.-B.: Single-image hdr reconstruction by learning to reverse the camera pipeline. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (2020)*, pp. 1651–1660. 2, 5
- [LLS*24] LEE B., LEE H., SUN X., ALI U., PARK E.: Deblurring 3d gaussian splatting, 2024. [arXiv:2401.00834](https://arxiv.org/abs/2401.00834). 2, 3, 6
- [LLSL23] LEE D., LEE M., SHIN C., LEE S.: DP-NeRF: Deblurred neural radiance field with physical scene priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023)*, pp. 12386–12396. 2, 3
- [LWZL22] LIU Z., WANG Y., ZENG B., LIU S.: Ghost-free high dynamic range imaging with context-aware transformer. In *European Conference on Computer Vision (2022)*, Springer, pp. 344–360. 3
- [MA21] MANTIUK R. K., AZIMI M.: Pu21: A novel perceptually uniform encoding for adapting existing quality metrics for hdr. In *2021 Picture Coding Symposium (PCS) (2021)*, IEEE, pp. 1–5. 6
- [MBRHD18] MARNERIDES D., BASHFORD-ROGERS T., HATCHETT J., DEBATTISTA K.: Expandnet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content. In *Computer Graphics Forum (2018)*, vol. 37, Wiley Online Library, pp. 37–49. 2
- [MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)* 41, 4 (2022), 1–15. 3

- [MHH23] MANTIUK R. K., HAMMOU D., HANJI P.: Hdr-vdp-3: A multi-metric for predicting image differences, quality and contrast distortions in high dynamic range and regular content. *arXiv preprint arXiv:2304.13625* (2023). 6
- [MHMB*22] MILDENHALL B., HEDMAN P., MARTIN-BRUALLA R., SRINIVASAN P. P., BARRON J. T.: Nerf in the dark: High dynamic range view synthesis from noisy raw images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 16190–16199. 2, 3, 8
- [MLL*22] MA L., LI X., LIAO J., ZHANG Q., WANG X., WANG J., SANDER P. V.: Deblur-nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 12861–12870. 2, 3, 6
- [MMB12] MITTAL A., MOORTHY A. K., BOVIK A. C.: No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing* 21, 12 (2012), 4695–4708. 6
- [MSB12] MITTAL A., SOUNDARARAJAN R., BOVIK A. C.: Making a “completely blind” image quality analyzer. *IEEE Signal processing letters* 20, 3 (2012), 209–212. 6
- [MST*21] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHY R., NG R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* 65, 1 (2021), 99–106. 1, 3
- [NWL*21] NIU Y., WU J., LIU W., GUO W., LAU R. W.: Hdr-gan: Hdr image reconstruction from multi-exposed ldr images with large motions. *IEEE Transactions on Image Processing* 30 (2021), 3885–3896. 3
- [PBM*22] PIDHORSKYI S., BAGAUTDINOV T., MA S., SARAGIH J., SCHWARTZ G., SHEIKH Y., SIMON T.: Depth of field aware differentiable rendering. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–18. 2
- [PC81] POTMESIL M., CHAKRAVARTY I.: A lens and aperture camera model for synthetic image generation. *ACM SIGGRAPH Computer Graphics* 15, 3 (1981), 297–305. 4, 10
- [PC82] POTMESIL M., CHAKRAVARTY I.: Synthetic image generation with a lens and aperture camera model. *ACM Transactions on Graphics (TOG)* 1, 2 (1982), 85–108. 2
- [PSA*21] PRABHAKAR K. R., SENTHIL G., AGRAWAL S., BABU R. V., GORTHI R. K. S. S.: Labeled from unlabeled: Exploiting unlabeled data for few-shot deep hdr deghosting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2021), pp. 4875–4885. 3
- [RFS22] RÜCKERT D., FRANKE L., STAMMINGER M.: Adop: Approximate differentiable one-pixel point rendering. *ACM Transactions on Graphics (ToG)* 41, 4 (2022), 1–14. 2
- [RHD*10] REINHARD E., HEIDRICH W., DEBEVEC P., PATTANAİK S., WARD G., MYSZKOWSKI K.: *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann Series in Computer Graphics. Elsevier Science, 2010. URL: https://books.google.de/books?id=wli_lkejoYcC.2
- [Sco92] SCOFIELD C.: 212-d depth-of-field simulation for computer animation. In *Graphics Gems III (IBM Version)*. Elsevier, 1992, pp. 36–38. 2
- [SF16] SCHONBERGER J. L., FRAHM J.-M.: Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 4104–4113. 4
- [SGHS98] SHADE J., GORTLER S., HE L.-W., SZELISKI R.: Layered depth images. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (1998), pp. 231–242. 3, 8
- [SKY*12] SEN P., KALANTARI N. K., YAESOUBI M., DARABI S., GOLDMAN D. B., SHECHTMAN E.: Robust patch-based hdr reconstruction of dynamic scenes. *ACM Trans. Graph.* 31, 6 (2012), 203–1. 2
- [SRK20] SANTOS M. S., REN T. I., KALANTARI N. K.: Single image hdr reconstruction using a cnn with masked features and perceptual loss. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 80–1. 2
- [SZW*23] SI H., ZHAO B., WANG D., GAO Y., CHEN M., WANG Z., LI X.: Fully self-supervised depth estimation from defocus clue. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 9140–9149. 2, 10
- [TWZ*23] TEL S., WU Z., ZHANG Y., HEYRMAN B., DEMONCEAUX C., TIMOFTE R., GINHAC D.: Alignment-free hdr deghosting with semantics consistent transformer. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)* (2023), IEEE, pp. 12790–12799. 3
- [VHM*22] VERBIN D., HEDMAN P., MILDENHALL B., ZICKLER T., BARRON J. T., SRINIVASAN P. P.: Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), IEEE, pp. 5481–5490. 3
- [WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612. 5, 6
- [WCS*22] WANG C., CHEN B., SEIDEL H.-P., MYSZKOWSKI K., SERRANO A.: Learning a self-supervised tone mapping operator via feature contrast masking loss. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 71–84. 5
- [WJ22] WON C., JEON H.-G.: Learning depth from focus in the wild. In *European Conference on Computer Vision* (2022), Springer, pp. 1–18. 2
- [WLP*22] WU Z., LI X., PENG J., LU H., CAO Z., ZHONG W.: DoF-NeRF: Depth-of-field meets neural radiance fields. In *Proceedings of the 30th ACM International Conference on Multimedia* (2022), pp. 1718–1729. 2, 3
- [WSP*23a] WANG C., SERRANO A., PAN X., CHEN B., MYSZKOWSKI K., SEIDEL H.-P., THEOBALT C., LEIMKÜHLER T.: Glowgan: Unsupervised learning of hdr images from ldr images in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2023), pp. 10509–10519. 2
- [WSP*23b] WANG C., SERRANO A., PAN X., WOLSKI K., CHEN B., MYSZKOWSKI K., SEIDEL H.-P., THEOBALT C., LEIMKÜHLER T.: An implicit neural representation for the image stack: Depth, all in focus, and high dynamic range. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–11. 2, 3, 10
- [XC21] XIONG P., CHEN Y.: Hierarchical fusion for practical ghost-free high dynamic range imaging. In *Proceedings of the 29th ACM International Conference on Multimedia* (2021), pp. 4025–4033. 3
- [YCH*24] YU Z., CHEN A., HUANG B., SATTLER T., GEIGER A.: Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 19447–19456. 3
- [YGS*19] YAN Q., GONG D., SHI Q., HENGEL A. V. D., SHEN C., REID I., ZHANG Y.: Attention-guided network for ghost-free high dynamic range imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 1751–1760. 3
- [YHS*23] YAN Q., HU T., SUN Y., TANG H., ZHU Y., DONG W., VAN GOOL L., ZHANG Y.: Towards high-quality hdr deghosting with conditional diffusion models. *IEEE Transactions on Circuits and Systems for Video Technology* (2023). 3
- [YLCL24] YAN Z., LOW W. F., CHEN Y., LEE G. H.: Multi-scale 3d gaussian splatting for anti-aliased rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 20923–20931. 3
- [YZC*22] YAN Q., ZHANG S., CHEN W., LIU Y., ZHANG Z., ZHANG Y., SHI J. Q., GONG D.: A lightweight network for high dynamic range imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 824–832. 3

- [YZC*23] YAN Q., ZHANG S., CHEN W., TANG H., ZHU Y., SUN J., VAN GOOL L., ZHANG Y.: Smae: Few-shot learning for hdr deghosting with saturation-aware masked autoencoders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 5775–5784. [3](#)
- [YZL*20] YAN Q., ZHANG L., LIU Y., ZHU Y., SUN J., SHI Q., ZHANG Y.: Deep hdr imaging via a non-local network. *IEEE Transactions on Image Processing* 29 (2020), 4308–4322. [3](#)
- [ZIE*18] ZHANG R., ISOLA P., EFROS A. A., SHECHTMAN E., WANG O.: The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR* (2018). [6](#)
- [ZMN*19] ZHANG X., MATZEN K., NGUYEN V., YAO D., ZHANG Y., NG R.: Synthetic defocus and look-ahead autofocus for casual videography. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–16. [3](#)
- [ZPVBG01] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: Ewa volume splatting. In *Proceedings Visualization, 2001. VIS'01.* (2001), IEEE, pp. 29–538. [3](#)
- [ZWL*24] ZHANG Z., WANG H., LIU S., WANG X., LEI L., ZUO W.: Self-supervised high dynamic range imaging with multi-exposure images in dynamic scenes. In *ICLR* (2024). [3](#)