





Dynamic Neural Radiance Field From Defocused Monocular Video

Xianrui Luo¹, Huiqiang Sun¹, Juewen Peng², and Zhiguo Cao¹

¹ Key Laboratory of Image Processing and Intelligent Control, Ministry of Education,
School of AIA, Huazhong University of Science and Technology, China
{xianruiluo,shq1031,zgcao}@hust.edu.cn

² College of Computing and Data Science, Nanyang Technological University,
Singapore
juewen.peng@ntu.edu.sg

<https://github.com/xianrui-luo/D2RF>

Abstract. Dynamic Neural Radiance Field (NeRF) from monocular videos has recently been explored for space-time novel view synthesis and achieved excellent results. However, defocus blur caused by depth variation often occurs in video capture, compromising the quality of dynamic reconstruction because the lack of sharp details interferes with modeling temporal consistency between input views. To tackle this issue, we propose D^2RF , the first dynamic NeRF method designed to restore sharp novel views from defocused monocular videos. We introduce layered Depth-of-Field (DoF) volume rendering to model the defocus blur and reconstruct a sharp NeRF supervised by defocused views. The blur model is inspired by the connection between DoF rendering and volume rendering. The opacity in volume rendering aligns with the layer visibility in DoF rendering. To execute the blurring, we modify the layered blur kernel to the ray-based kernel and employ an optimized sparse kernel to gather the input rays efficiently and render the optimized rays with our layered DoF volume rendering. We synthesize a dataset with defocused dynamic scenes for our task, and extensive experiments on our dataset show that our method outperforms existing approaches in synthesizing all-in-focus novel views from defocus blur while maintaining spatial-temporal consistency in the scene.

Keywords: Dynamic Novel View Synthesis · Neural Radiance Field · Depth-of-Field

1 Introduction

Capturing videos from cameras or smartphones has become a new norm for daily life. However, videos are typically recorded from a monocular camera, restricting the captured scene to fixed viewpoints. To depict the dynamic scene from flexible viewpoints, dynamic view synthesis [20, 21, 34] is proposed to facilitate the generation of photorealistic novel views from arbitrary camera angles and perspectives, thus allowing for free viewpoints. The ability to capture a 3D dynamic scene in




Fig. 1: Given a monocular video captured with defocus blur, existing dynamic NeRF approaches fail to recover high-quality details and tend to produce blurry views, and our method D^2RF synthesizes sharp novel views.

a monocular setting enhances the importance of dynamic novel view synthesis in video-related tasks such as video stabilization, augmented reality, and view interpolation. To interpolate new views from existing sparse input views, recent works use neural networks to train a neural radiance field (NeRF) [27]. The network is a multilayer perceptron (MLP) that learns to map spatial coordinates and view directions to volume densities and view-dependent RGB colors.

Recent dynamic NeRF methods [20, 21, 31] have made progress by establishing spatial-temporal consistency to achieve space-time novel view synthesis. However, these methods are designed under the assumption that the model is trained on all-in-focus image sequences, where all objects within the scene remain focused and sharp during the entire filming. Defocus blur has not been addressed by current dynamic NeRF models, and the blur can result in performance degradation for the reconstructed dynamic NeRF, as it poses challenges for modeling dynamic object motions due to the lack of sharp details and the inability to model temporal consistency. As shown in Fig. 1, current dynamic NeRF methods perform poorly when applied to monocular videos affected by defocus blur, resulting in an inability to restore sharp contents. Maintaining all objects in the scene in focus may offer a solution, however, maintaining every frame in focus is challenging when capturing monocular videos. Defocus blur occurs when the camera’s aperture is wide, and the objects are not positioned at the focal distance. Therefore, defocus blur tends to manifest in videos due to the scene comprising multiple objects with significant depth variations coupled with the common equipment of large aperture for more natural light and enhanced imaging quality. Furthermore, the focal distance may constantly fluctuate in frames if the photographer lacks expertise. Considering that defocus blur is almost unavoidable during video capture, recovering a sharp dynamic NeRF is a valuable problem yet to be explored.

Several works have been proposed to address defocus blur from static scenes. Physical scene prior and kernel estimation [16, 25] are introduced to mitigate

blurring, and an explicit scatter-gather technique [50] is proposed to better model the defocus blur. However, these methods are specifically tailored for static multi-view inputs and are not capable of representing defocused dynamic scenes. Previous methods apply blurring after the last step of NeRF: volume rendering. The blurring is before optimization but after the typical rendering.

To tackle defocus blur in monocular dynamic videos, we propose D^2RF , a dynamic framework designed to incorporate the defocus blur model into dynamic NeRF training, enabling the recovery of sharp novel views from defocused monocular videos. Our approach involves modeling defocus blur by estimating blur kernels and introducing layered Depth-of-Field (DoF) volume rendering to the input rays. We connect the classical volume rendering with DoF rendering, seamlessly integrating the blurring process into the NeRF pipeline. We modify the layer-based kernel in DoF rendering to the ray-based kernel to fit into volume rendering in NeRF. To decrease computational costs, we transition from using a traditional disk blur kernel to employing a sparse one [25] and optimize the rays originating from these sparse points. We use an MLP to predict both the 3D positions of the kernel points and their corresponding weights. Subsequently, the optimized rays and the estimated kernels serve as inputs for the aforementioned layered DoF volume rendering process. It is worth noting that our D^2RF enables the dynamic NeRF model to learn a sharp scene representation from defocused input views, leading to sharp novel views.

To improve the quality of reconstructed static regions, we utilize the blending technique [20] to independently predict the static and dynamic scenes. The blended color is then rendered with our layered DoF volume rendering pipeline. Furthermore, we adopt a cross-time rendering approach to represent temporal consistency in dynamic scenes. This approach models scene correlation across multiple input views and utilizes deformed rays for layered DoF volume rendering. Integrating scene information from other frames into the target frame enhances the representation of dynamic scenes.

To address the absence of existing dynamic NeRF defocus deblurring datasets, we collect and synthesize defocused video sequences from a stereo dataset [49] and conduct an extensive analysis to demonstrate D^2RF can effectively address defocus blur inputs and render high-quality sharp novel views. We compare with current dynamic NeRF methods as well as the combination of single-image deblurring and dynamic NeRF. Results show that D^2RF outperforms existing methods. Additionally, we conduct ablation studies to validate the effectiveness of each component. In summary, our main contributions are as follows.

We present the first dynamic NeRF model D^2RF designed to recover sharp novel views from a defocused monocular video.

We introduce layered DoF volume rendering in dynamic NeRF to model the defocus blur and restore a sharp representation.

We transform the layered blur kernel to the ray-based kernel to fit into NeRF training and incorporate the optimized sparse kernel with the layered DoF volume rendering, enabling an efficient blurring process.

2 Related Work

Depth-of-Field Rendering. Depth-of-field (DoF) rendering, also known as bokeh rendering, generates realistic defocus blur in out-of-focus areas to enhance the visual prominence of the focused subject within an image. Physically based methods [1, 18] rely on 3D scene information and they are time-consuming. Recent methods use neural networks [10, 11] for end-to-end training. DeepFocus [52] entails an exact depth map to render realistic bokeh. However, the exact depth map is hard to acquire. Therefore, some methods [45, 48] instead predict depth for DoF rendering. To automate the rendering process, bokeh is directly rendered [9, 24] from an all-in-focus input without depth maps and controlling parameters. Synthesized data [32] from ray tracing is introduced for training to combine classical and neural rendering. To generate partial occlusion, some methods [33, 38] decompose the scene into multiple layers and separately blur each layer before compositing them, inspired by layered rendering [4, 55]. In this work, we show the connection between DoF rendering and NeRF rendering, integrating the layered blurring kernels with volume rendering.

Image Defocus Deblurring. To recover a sharp image from defocus blur, traditional methods conduct a two-stage pipeline: (1) estimate a defocus map [29, 39] from image priors, (2) the defocus map is used as guidance to deblur the image from non-blind deconvolution [6, 19]. With the current advances in neural networks, recent methods train end-to-end on datasets [2]. Modifications on convolutions [17, 42] are proposed to improve the deblurring performance. The Transformer architecture and a novel multi-scale feature extraction module [53] are introduced for high-resolution restoration. Additional inputs, such as dual pixel [2], light field [36], video sequences [3], and depth images [28] are utilized to guide defocus deblurring. Current methods do not take multi-view 3D geometry and dynamic scenes into account. We aim to explore defocus deblurring in dynamic NeRF synthesis, and our method integrates the DoF rendering into dynamic NeRF to synthesize defocus blur, which in turn facilitates deblurring.

Neural Radiance Field. Neural radiance field (NeRF) is an implicit representation [23, 26, 40, 41] for photo-realistic novel view synthesis, which aims to generate new camera perspectives given a specified set of input viewpoints. Different from explicit representations [7, 12], NeRF [27] learns an implicit continuous function to model the complex geometry and appearance of a scene. Subsequent studies aim to synthesize high-quality novel views from abnormal inputs. To recover a sharp NeRF from defocus-blurred inputs, Deblur-NeRF [25] proposes a deformable sparse kernel to model defocus blur, and DP-NeRF [16] adopts physical scene priors to mitigate blurring. DoF-NeRF [50] introduces the Concentrate-and-Scatter technique to explicitly enable controllable DoF effects. However, all these methods render DoF in a post-process manner and do not consider the volume sampling process. LensNeRF [14] restores a sharp NeRF from defocused static scene. The technique is valid for static scenes but does not model the temporal photometric consistency between video frames. Current methods are only effective in static scenes, where the objects are not moving. Dynamic NeRF synthesis with defocus blur remains unexplored.

Dynamic Neural Radiance Field. Recent works extend NeRF from static scenes to dynamic scenes [21, 51]. Given a monocular video, some approaches apply a deformation field on a canonical representation [30, 31, 34, 44, 47] for space-time novel view synthesis. Scene flow also works as a constraint to establish temporal consistency and model object motions in dynamic scenes [8, 20, 22, 46]. These methods require sharp inputs and fail when the defocus blur is in the input monocular video. Compared with existing methods, we synthesize novel views given the defocus blur. Furthermore, we propose to tackle the blur by introducing the layered DoF volume rendering in dynamic NeRF rendering.

3 Method

We propose D^2RF , a dynamic radiance field pipeline to recover sharp novel views from a defocused monocular video. In the following, we first briefly revisit the concept of NeRF and the principle of DoF rendering. In section 3.1, we analyze the connection and the distinction between DoF rendering and volume rendering in NeRF training, proposing layered DoF volume rendering with optimized ray-based blur kernels. Then we show how to represent dynamic scenes in a space-time manner, blend dynamic and static scenes with layered DoF volume rendering, and ensure temporal consistency within the scene in Section 3.2. Finally, we provide the training details in Section 3.3.

Preliminary: Neural Radiance Field Our method follows the principle of Neural Radiance Field (NeRF), which represents a static 3D scene as an implicit continuous function F_Θ . This function is parameterized by a multilayer perceptron (MLP) and it takes a spatial point location $\mathbf{x} \in \mathbb{R}^3$ and viewing direction $\mathbf{d} \in \mathbb{R}^3$ as inputs. F_Θ maps the inputs to color \mathbf{c} and volume density σ :

$$(\mathbf{c}, \sigma) = F_\Theta(\mathbf{x}, \mathbf{d}). \quad (1)$$

To acquire the RGB value of a pixel coordinate p , NeRF first specifies a ray $\mathbf{r}_p(t) = \mathbf{o} + t\mathbf{d}_p$ emitted from camera projection center \mathbf{o} along the viewing direction \mathbf{d}_p . Then followed by the classical volume rendering technique [13], the color of $\mathbf{r}_p(t)$ can be represented as an integral of all colors along the ray, computed as:

$$\hat{C}(\mathbf{r}_p) = \sum_{i=1}^k T_i (1 - \exp(-\sigma \delta_i)) \mathbf{c}(\mathbf{r}_p(t_i), \mathbf{d}_p), \quad (2)$$

where

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

is the accumulated transmittance along ray \mathbf{r}_p , and $\delta_i = t_{i+1} - t_i$ is the distance between the two sampled points. The vanilla NeRF assumes a pinhole camera model, where the color of a pixel $\hat{C}(\mathbf{r}_p)$ is calculated by points on a single ray $\mathbf{r}_p(t)$, so it fails to model the defocus blur.




Fig. 2: The defocus blur formation. We model the defocus blur of the center purple pixel. The blur and green objects are from different layers. We define the light path as the left-right view and the image plane shows the front-back view of the defocus blur. The CoC of the purple point gathers its neighboring pixels on the green foreground and those on the blue background. This defocus modeling indicates that the object originally occluded (red square from the blue object) by the green object under the pinhole camera view can contribute to the rendered purple point by the green and blue semi-circles. In Eq.5-7 we model the defocus blur from layer visibility W_i , then use the link between visibility and opacity to integrate DoF and volume rendering in Eq.8. The visibility is learned from the layered DoF volume rendering pipeline.

Preliminary: Defocus Blur and DoF Rendering To model the defocus blur and render DoF, first we introduce bokeh rendering. Bokeh is the aesthetic quality of the blur in the out-of-focus region, caused by Circle of Confusion (CoC). Previous works [32, 45] have shown the CoC formation process. The radius γ of each pixel is

$$\gamma = af \left| \frac{1}{z} - \frac{1}{z_f} \right|, \quad (4)$$

where a is the camera aperture radius, f is the focal length. z is the depth of the pixel and z_f is the focal distance. Now that we have the blur radius of each pixel, we adopt layered bokeh rendering [33]. To blur the input image with kernels, classical bokeh rendering can be categorized into *gather* and *scatter* operations. We chose *gather* because it fits the layered bokeh rendering pipeline and is compatible with the volume rendering in NeRF described in section 3.1. We define the layer as a set of fronto-parallel planes at fixed depths by discretization. Now we have the blur kernel $K(\gamma_i)$ for layer i , we apply an alpha composition for DoF rendering. The rendered DoF image B is formulated as:

$$B = \frac{\psi(C_i)}{\psi(1)}, \quad (5)$$

where

$$\psi(X) = \sum_{i=1}^N \prod_{j=i+1}^N (1 - W_j - K(\gamma_j))(XW_i - K(\gamma_i)). \quad (6)$$

$\psi(X)$ is the compositing blurring function with the image X as input, ψ is the convolution operation, N is the number of layers, 1 is an image with all values

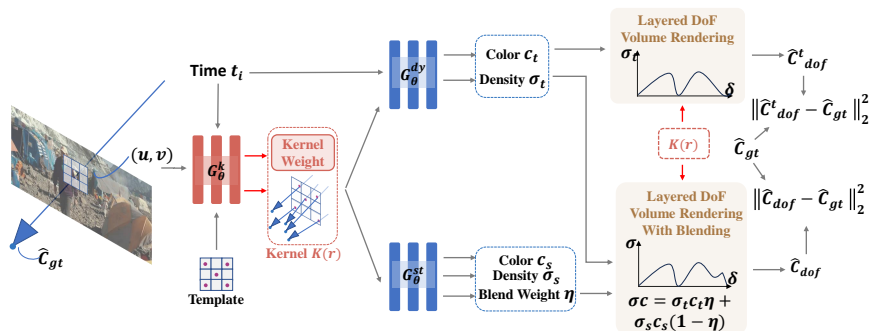


Fig. 3: Pipeline of our framework. The framework takes a set of plane coordinates (u, v) , the time embedding t_i , and defines a kernel template as inputs, the outputs are the blur kernels $K(\mathbf{r}_i)$ consisting of the sparse optimized rays with their corresponding weights. The rays are then fed to the two MLPs G_θ^{st} and G_θ^{dy} to independently represent static and dynamic scenes. The final color is rendered by layered DoF volume rendering (Section 3.1), $\hat{C}_{dof}(\mathbf{r})$ is from Eq.8 and $\hat{C}_{dof}^t(\mathbf{r})$ is from Eq.13. The rendered defocused results (dynamic and blended) are supervised by the input defocused views. For testing we directly render the rays without layered DoF volume rendering and the kernel.

as 1. Each layer i encodes an RGB image C_i , with their corresponding visibility weight W_i , ensuring the correct blending of different layers. We show the defocus blur formation and how rays gather between layers in Fig. 2, demonstrating why we should apply the visibility term in DoF rendering. This layered visibility-aware rendering gathers the rays with inter-layer weights for each layer, then blends the layers with the weights from back to front.

3.1 Layered DoF Volume Rendering

Now that we have a basic understanding of NeRF and DoF rendering, we introduce our layered DoF volume rendering module. As shown in Fig. 3, we introduce layered DoF volume rendering, incorporating DoF rendering into volume rendering in dynamic NeRF. To seamlessly fuse Eq. 5 and Eq. 6 in NeRF training, we utilize the connections between the blur formation and volume rendering in Eq. 2. The volume rendering assumes that, for each ray $\mathbf{r}(t_i)$ from a scene, the sampling points in the volume all have an alpha value $\alpha_i = 1 - \exp(-\sigma\delta_i)$, which indicates the opacity of the point. A larger alpha value shows greater absorption of light, making it more likely for the light to pass through that point. In contrast, a smaller alpha value means less absorption of light at that point, and the object at that point is more transparent. In Eq. 6, the term W_i represents the visibility between the current layer i and all the layers behind it. This visibility W has the same physical meaning as the opacity value α in volume rendering. Furthermore, in NeRF rendering, the volume of the scene is sampled in a discretized manner, which typically involves uniformly sampling discrete points along the ray path. This layer discretization is compatible with the one in DoF rendering.

Therefore we first transform Eq. 5 from the format of layered images into the format of individual layered pixels:

$$B(\mathbf{r}) = \frac{\sum_{i=1}^k \left(\prod_{j=i+1}^k (1 - w_j K(\mathbf{r})) w_i \mathbf{c}_i K(\mathbf{r}) \right)}{\sum_{i=1}^k \left(\prod_{j=i+1}^k (1 - w_j K(\mathbf{r})) w_i K(\mathbf{r}) \right)}, \quad (7)$$

where \mathbf{c}_i and w_i are the colors and layer visibilities of the pixels surrounding \mathbf{r} within the blur kernel $K(\mathbf{r})$. Then we propose to integrate layered DoF volume rendering into dynamic NeRF training, the rendering can be formulated as:

$$\hat{C}_{doF}(\mathbf{r}_p) = \frac{\sum_{i=1}^k \left((T_i K(\mathbf{r})) (1 - \exp(-\sigma \delta_i)) \mathbf{c}(\mathbf{r}(t_i), \mathbf{d}) K(\mathbf{r}) \right)}{\sum_{i=1}^k \left((T_i K(\mathbf{r})) (1 - \exp(-\sigma \delta_i)) K(\mathbf{r}) \right)}, \quad (8)$$

where the variable name is consistent as Eq. 2 and Eq. 7, *e.g.* k is the number of sampling layers both in DoF rendering and volume rendering, \mathbf{c}_i is the emitted color of layer i blurred by $K(\mathbf{r})$.

Although DoF rendering and NeRF volume rendering share similar traits, there is a difference in their implementations, as the DoF rendering is originally used on different layers of the entire image, and NeRF training is optimized by rays sampled across the near and far planes. To implement the DoF rendering in Eq. 8, we switch the layer-based kernels $K(\gamma_i)$ to ray-based kernels $K(\mathbf{r})$, ensuring that the blur kernels are still diverse based on different positions. $K(\gamma_i)$ is defined by the divided layers, and $K(\mathbf{r})$ is defined by the input rays. Furthermore, to alleviate the high computational cost, we apply a sparse kernel following [25], which replaces the dense kernel with a smaller number of sparse points. The convolution process with $K(\mathbf{r}_p)$ can be formulated as:

$$\mathbf{b}_p = \sum_{j \in S(p)} \mathbf{c}_j g_j, \quad (9)$$

where \mathbf{b}_p is the result of point p after convolution, $S(p)$ represents the sparse points surrounding point p , g_j and \mathbf{c}_j are the corresponding weights and colors from the sparse points. We set $\sum_{j \in S(p)} g_j = 1$ to ensure the brightness consistency. Though we fix the number of sparse points, we additionally apply ray deformation as [25] to optimize the rays as the spatially varying real-world blur kernels. As shown in Fig. 3, we jointly optimize the translation of the center ray origin of the blur kernel. For a plane coordinate at the center of kernel (u, v) , we use an MLP G_θ^k to predict the deformation of the kernel points as well as the kernel weights:

$$(\Delta \mathbf{j}, g_j) = G_\theta^k((u, v), \mathbf{j}, t_l), \quad (10)$$

where t_l is the embedding of the input frame, \mathbf{j} are the original rays from the kernel template, g_j are the corresponding kernel weights, and the final optimized ray inputs are calculated as $\mathbf{r}_j = \mathbf{j} + \Delta \mathbf{j}$.

3.2 Dynamic Radiance Field

With the layered DoF volume rendering module in place, we turn to the dynamic scene representation. Given a monocular video of a dynamic scene, we take the video frames and camera parameters as inputs. Following previous dynamic NeRF methods, we represent the scene with an MLP G_θ^{dy} , which maps the input spatial point location \mathbf{x} , viewing direction \mathbf{d} , and time t , to the corresponding color, volume density, scene flow, and disocclusion weight:

$$(\mathbf{c}_t, \sigma_t, f_t, W_t) = G_\theta^{\text{dy}}(\mathbf{x}, \mathbf{d}, t), \quad (11)$$

where G_θ^{dy} is parameterized by θ . The scene flow f_t and the disocclusion weight W_t are used for the following cross-time rendering. Based on the representation, we next describe our twist on the dynamic radiance field.

Blending Dynamic and Static Scene. We use two MLPs to model static and dynamic scenes as shown in Fig.3. One MLP is designed for modeling moving objects by various loss terms (cross-time *etc.*), and the other learns the static scene, allowing for more stable rendered still objects. To improve the quality of time-invariant static regions, we apply a blending technique as previous dynamic methods [20, 22], using an additional MLP to represent the static scene:

$$(\mathbf{c}, \sigma, \eta) = G_\theta^{\text{st}}(\mathbf{x}, \mathbf{d}), \quad (12)$$

where η denotes the blending weight used to blend static and dynamic models. We denote the static color as \mathbf{c}_s , the dynamic color as \mathbf{c}_t , the static volume density and the dynamic volume density as σ_s and σ_t . As shown in Fig. 3, the layered DoF volume rendering with the blending weight is defined as:

$$\hat{C}_{\text{dof}}(\mathbf{r}) = \frac{\sum_{i=1}^k T_i^b \text{volume}_i^\epsilon}{\sum_{i=1}^k T_i^b \text{volume}_i^\epsilon}, \quad (13)$$

where $T_i^b = \exp(-\sum_{j=1}^{i-1} \sigma_s \sigma_t \delta_j) K(\mathbf{r})$. volume_i^ϵ and volume_i^ϵ are rendered with $\eta(t)$ as weights:

$$\begin{aligned} \text{volume}_i^\epsilon &= \eta(t) (1 - \exp(-\sigma_s \delta_i)) \mathbf{c}_s(\mathbf{r}(t), \mathbf{d}) K(\mathbf{r}) \\ &+ (1 - \eta(t)) (1 - \exp(-\sigma_t \delta)) \mathbf{c}_t(\mathbf{r}(t), \mathbf{d}) K(\mathbf{r}), \end{aligned} \quad (14)$$

and

$$\begin{aligned} \text{volume}_i^\epsilon &= \eta(t) (1 - \exp(-\sigma_s \delta_i)) K(\mathbf{r}) \\ &+ (1 - \eta(t)) (1 - \exp(-\sigma_t \delta)) K(\mathbf{r}). \end{aligned} \quad (15)$$

Finally, we apply the rendering loss for the blending results:

$$L_{\text{color}}^b = k \hat{C}_{\text{dof}}(\mathbf{r}) - \hat{C}_{\text{gt}}(\mathbf{r}) k_2^2. \quad (16)$$

Apart from blending the static and dynamic scenes, we also constrain the rendered result of dynamic scene $\hat{C}_{\text{dof}}^t(\mathbf{r})$ by $L_{\text{color}}^t = k \hat{C}_{\text{dof}}^t(\mathbf{r}) - \hat{C}_{\text{gt}}(\mathbf{r}) k_2^2$.

Cross-time rendering. Suppose we only apply photometric consistency on individual time stamps like static NeRF. In that case, the model might fail to build temporal consistency in dynamic scenes, resulting in overfitting input views and failing to synthesize correct novel views. Therefore, temporal consistency between neighboring frames [8, 21] is crucial for training dynamic radiance fields. To establish temporal consistency between corresponding frames, we apply cross-time rendering under layered DoF volume rendering. For a timestamp t_m from the m -th frame, we assign the three-dimensional scene flow of a spatial point \mathbf{x} from t_m to timestamp t_n as $f_{\mathbf{x}}^m(n)$, where n denotes the neighbor frames of m . The corresponding point of \mathbf{x} at t_n can be computed as $\mathbf{x}^{m \rightarrow n} = \mathbf{x} + f_{\mathbf{x}}^m(n)$. This correspondence works on all the sampled points of a ray, therefore, we use the layered DoF volume rendering to acquire the color of the warped ray:

$$\hat{C}_{dof}^{m \rightarrow n}(\mathbf{r}) = \frac{\sum_{i=1}^k \left((T_i^n \cdot K(\mathbf{r})) (1 - \exp(-\sigma_n \delta_i)) \mathbf{c}_t(\mathbf{r}(t_i)^{m \rightarrow n}, \mathbf{d}) \cdot K(\mathbf{r}) \right)}{\sum_{i=1}^k \left((T_i^n \cdot K(\mathbf{r})) (1 - \exp(-\sigma_n \delta_i)) \cdot K(\mathbf{r}) \right)}, \quad (17)$$

where $T_i^n = \exp(-\sum_{j=1}^i \sigma_n \delta_j)$, and $\mathbf{r}(t_i)^{m \rightarrow n} = \mathbf{r}(t_i)^m + f_{\mathbf{p}}^m(n)$. Then we compare the rendering loss between $\hat{C}_{dof}^{m \rightarrow n}(\mathbf{r})$ and $\hat{C}_{gt}^m(\mathbf{r})$:

$$\mathcal{L}_{cross} = \sum_{n \in \mathcal{N}(m)} \mathbf{W}^{n \rightarrow m}(\mathbf{r}) k \hat{C}_{dof}^{m \rightarrow n}(\mathbf{r}) - \hat{C}_{gt}^m(\mathbf{r}) k_2^2, \quad (18)$$

where $n \in \mathcal{N}(m)$ are the neighboring frames of the m -th frame, $\mathbf{W}^{n \rightarrow m}(\mathbf{r})$ is calculated from volume rendering the weight W_t in Eq. 11. The principle of disocclusion weight comes from the ambiguity of scene flow. Although scene flow is effective in establishing pixel-level alignment, it performs poorly when the pixels in one image are occluded. Therefore, we need a confidence map to demonstrate whether the pixel is occluded and the flow is correct. A possible solution is to use forward-backward consistency check [5] to alleviate the problem. Here we let the network learn the weight \mathbf{W} in an unsupervised manner [20].

3.3 Optimization

Monocular reconstruction of complex dynamic scenes is highly ill-posed and prone to local minima during optimization [8, 20]. Therefore, we introduce data-driven priors \mathcal{L}_{data} to supervise the scene geometry apart from the previous photometric constraint. Similar to previous works [20], \mathcal{L}_{data} consists of (1) scale-shift invariant monocular depth loss and (2) scene flow consistency and L1 regularization. We use RAFT [43] and DPT [35] to obtain ground truth optical flow and depth for the input images.

Implementations. We set the max kernel size as 10, which denotes the sparse kernel points are inquired under a radius of 10. The number of kernel points is 5. We employ the Adam optimizer [15] to jointly optimize the static and

dynamic MLPs and the blur kernels. We use COLMAP [37] to estimate the camera intrinsics and extrinsics. The learning rate is $5 \cdot 10^{-4}$. We train each scene for 250k iterations, which takes around two days on a single NVIDIA RTX 3090 GPU. The rendering takes roughly 13 seconds for each 940 × 360 frame.

4 Experiments

There are no available datasets for defocused inputs in dynamic NeRF, so we curated 8 defocused dynamic scenes from an existing dataset VDW [49] for evaluation. VDW consists of sharp stereo image sequences and their corresponding disparity sequences. We then use the bokeh rendering pipeline in BokehMe [32] for defocus blur generation. Given the disparity sequences, the synthesized bokeh quality surpasses current classical rendering and neural rendering methods. We collected 8 dynamic scenes of 1880 × 720 suitable for our task from this dataset, the image is resized to 940 × 360 for training. To closely resemble actual capturing scenarios, unlike static blur dataset [25] with a random focusing scheme, we gradually adjust the focal distance along the scene disparity, mimicking the typical focusing process observed in videos. The aperture is fixed for each scene. We utilize the left-view defocused image sequences for training and the corresponding right-view sharp image sequences for evaluation. The 8 scenes are named Camp, Shop, Car, Mountain, Dining1, Dining2, Dock, Gate.

4.1 Baseline Comparisons

The baselines include two types: 1) state-of-the-art Dynamic NeRF, *e.g.* scene flow methods NSFF [20] and DVS [8], deformation-based method HyperNeRF [31], robust dynamic method RoDynRF [22] and 2) an image-space baseline that uses single-view image deblurring [42] then trains the dynamic NeRF with the deblurred inputs. As all the current dynamic methods are not designed for defocus blur, we choose the robust RoDynRF and the flow-based DVS to train with [42] for baseline comparison. We use the pre-trained deblurring model for fairness.


4.2 Qualitative Results

We show qualitative comparisons in Fig. 4 and Fig. 5. In Fig. 4, we compare our method with existing dynamic NeRF methods, and in Fig. 5, we compare our approach with the chosen dynamic NeRF models and their counterparts with 2D image-deblurring for preprocessing. We present the results on all scenes in these two figures to show our method’s effectiveness, 4 scenes for each.

One can observe: 1) current dynamic NeRF methods can not recover sharp details from the defocus blur; 2) the single-view 2D deblurring preprocessing [42] can somewhat recover sharp details in some regions, however, each view with the deblurred content lacks consistency across views, and the performance is still less stable and reliable than our method; 3) our method D^2RF outperforms these baselines in handling defocus blur and generating sharp novel views. please refer to the supplementary material for more results.

Table 1: Quantitative comparisons against all baselines. The best performance is **boldfaced**, and the second is underlined.

Methods	PSNR	SSIM	LPIPS#	Methods	PSNR	SSIM	LPIPS#
HyperNeRF [31]	26.96	0.780	<u>0.208</u>	RoDynRF [22]	<u>26.18</u>	0.770	0.227
RoDynRF [22]	26.18	0.770	0.227	DVS [8]	25.43	0.764	0.242
NSFF [20]	<u>27.01</u>	<u>0.803</u>	0.209	[42] + [22]	25.79	<u>0.776</u>	<u>0.196</u>
DVS [8]	25.43	0.764	0.242	[42] + [8]	24.52	0.757	0.208
D^2RF (Ours)	27.30	0.816	0.130	D^2RF (Ours)	27.30	0.816	0.130

**Fig. 4: The qualitative results with all dynamic NeRF baselines.** Compared with existing dynamic NeRF methods, our method generates sharper novel views that are more faithful and have more details. The scenes are Mountain, Shop, Car, Dock.

4.3 Quantitative Results

We compare D^2RF with the 6 baseline methods quantitatively on our synthesized dataset. We use PSNR, SSIM, and LPIPS [54] as metrics for evaluation.

As shown in Table 1, D^2RF outperforms other state-of-the-art baselines. As mentioned above, to demonstrate the superiority of our method in modeling defocus blur in 3D space, we choose two dynamic NeRF baselines RoDynRF and DVS, and train these dynamic NeRF models while preprocessing the input blurry images with 2D defocus deblurring for comparison. This comparison highlights that, although the 2D deblurring method alleviates the defocus blur on each input view, this method struggles to maintain the consistency of scene information and tends to be unstable across all scenes. When dynamic NeRFs are trained on images deblurred with 2D methods, it results in undesirable artifacts




Fig. 5: The qualitative results with dynamic NeRF and their corresponding 2D image deblurring baselines. Although 2D image deblurring helps to alleviate the blur for dynamic NeRF in novel views (red box), our method is more stable and generates more reliable sharp details. The scenes are Camp, Dining1, Dining2, Gate.

in scene geometry and unsatisfying performance overall. In contrast, our method models defocus blur in 3D space across all input views, ensuring spatial-temporal consistency of scene geometry. We show the average results on all scenes, please refer to the supplementary material for separate results on all scenes.

4.4 Ablation Study

The main idea of D^2RF is to model the defocus blur by optimizing sparse blur kernels and incorporating the blur model into NeRF training. Therefore we conduct ablation studies on the layered DoF volume rendering module and the optimized kernels. Furthermore, we evaluate the effects of the static-dynamic blending scheme. As shown in Fig. 6 and Table 2, our framework works best when all the components are applied. We remove (1) the layered volume rendering module, and we use the kernel weight to directly gather the final rendered color as post-processing; (2) the optimized kernel, and we apply the full regular kernel and learn bokeh parameters; (3) static representation. We predict motion masks [20] to calculate the results on dynamic regions. We show that (1) the layered DoF volume rendering helps to model a more accurate defocus blur by integrating the blur into volume rendering, facilitating the recovery of sharp novel views; (2) the optimized kernel along with the layered DoF volume rendering provides an efficient and effective blurring pipeline for training; (3) the reconstruction of the static scene improves the performance because some

Table 2: Ablation study. The best performance is **boldfaced**, and the second is underlined. The left part calculates the results in dynamic regions, and the right part shows the results on the whole image.

Method (Dynamic)	PSNR	SSIM	LPIPS#	Method (All)	PSNR	SSIM	LPIPS#
w/o cross-time	19.46	0.607	0.351	w/o cross-time	22.61	0.725	0.232
w/o layered volume	24.42	<u>0.738</u>	<u>0.170</u>	w/o layered volume	27.11	<u>0.811</u>	0.211
w/o optimized kernel	<u>24.54</u>	0.735	0.246	w/o optimized kernel	<u>27.25</u>	0.795	0.216
w/o static	23.93	0.719	0.179	w/o static	26.20	0.769	<u>0.177</u>
Full (Ours)	24.60	0.747	0.162	Full (Ours)	27.30	0.816	0.130



Fig. 6: The visualizations of the ablation study. The corresponding error map is visualized at the bottom, where darker regions indicate smaller errors.

regions are static and an extra scene representation helps to stabilize training. We choose two scenes for visualization and show average quantitative results on all scenes, please refer to the supplementary material for more results.

5 Conclusion

We present D^2RF , a novel framework for all-in-focus dynamic novel view synthesis from defocused monocular videos. To tackle defocus blur from video capture, we integrate the blur model into NeRF training. Particularly, we connect DoF rendering with volume rendering, proposing layered DoF volume rendering for dynamic NeRF. We modify the layered kernel to the ray-based kernel and apply an optimized sparse kernel to gather the rays. We conduct experiments to show our method can recover from defocus blur and produce satisfying results. Although this framework works well overall, some limitations remain, such as the inability to handle extreme defocus blur. We plan to address these issues in future work. We declare to release the source code upon acceptance of the paper.

References

1. Abadie, G., McAuley, S., Golubev, E., Hill, S., Lagarde, S.: Advances in real-time rendering in games. In: ACM SIGGRAPH 2018 Courses, pp. 1–1 (2018) [4](#)
2. Abuolaim, A., Brown, M.S.: Defocus deblurring using dual-pixel data. In: Proceedings of the European conference on computer vision (ECCV). pp. 111–126. Springer (2020) [4](#)
3. Abuolaim, A., Delbracio, M., Kelly, D., Brown, M.S., Milanfar, P.: Learning to reduce defocus blur by realistically modeling dual-pixel data. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2289–2298 (2021) [4](#)
4. Busam, B., Hog, M., McDonagh, S., Slabaugh, G.: Sterefo: Efficient image refocusing with stereo vision. In: Proc. IEEE International Conference on Computer Vision Workshops (ICCVW). pp. 0–0 (2019) [4](#)
5. Chen, Q., Koltun, V.: Full flow: Optical flow estimation by global optimization over regular grids. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4706–4714 (2016) [10](#)
6. Dong, J., Roth, S., Schiele, B.: Dwdn: deep wiener deconvolution network for non-blind image deblurring. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(12), 9960–9976 (2021) [4](#)
7. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 605–613 (2017) [4](#)
8. Gao, C., Saraf, A., Kopf, J., Huang, J.B.: Dynamic view synthesis from dynamic monocular video. In: ICCV. pp. 5712–5721 (2021) [5](#), [10](#), [11](#), [12](#)
9. Ignatov, A., Patel, J., Timofte, R.: Rendering natural camera bokeh effect with deep learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 418–419 (2020) [4](#)
10. Ignatov, A., Patel, J., Timofte, R., Zheng, B., Ye, X., Huang, L., Tian, X., Dutta, S., Purohit, K., Kandula, P., et al.: Aim 2019 challenge on bokeh effect synthesis: Methods and results. In: Proc. IEEE International Conference on Computer Vision Workshops (ICCVW). pp. 3591–3598. IEEE (2019) [4](#)
11. Ignatov, A., Timofte, R., Qian, M., Qiao, C., Lin, J., Guo, Z., Li, C., Leng, C., Cheng, J., Peng, J., et al.: Aim 2020 challenge on rendering realistic bokeh. In: Proc. European Conference on Computer Vision Workshops (ECCVW). pp. 213–228. Springer (2020) [4](#)
12. Jimenez Rezende, D., Eslami, S., Mohamed, S., Battaglia, P., Jaderberg, M., Heess, N.: Unsupervised learning of 3d structure from images. *Advances in neural information processing systems* **29**, 4996–5004 (2016) [4](#)
13. Kajiya, J.T., Von Herzen, B.P.: Ray tracing volume densities. *ACM SIGGRAPH computer graphics* **18**(3), 165–174 (1984) [5](#)
14. Kim, M.J., Gu, G., Choo, J.: Lensnerf: Rethinking volume rendering based on thin-lens camera model. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3182–3191 (2024) [4](#)
15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [10](#)
16. 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 (CVPR). pp. 12386–12396 (June 2023) [2](#), [4](#)

17. Lee, J., Son, H., Rim, J., Cho, S., Lee, S.: Iterative filter adaptive network for single image defocus deblurring. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2034–2042 (2021) [4](#)
18. Lee, S., Eisemann, E., Seidel, H.P.: Real-time lens blur effects and focus control. *ACM Transactions on Graphics (TOG)* **29**(4), 1–7 (2010) [4](#)
19. Levin, A., Fergus, R., Durand, F., Freeman, W.T.: Image and depth from a conventional camera with a coded aperture. *ACM transactions on graphics (TOG)* **26**(3), 70–es (2007) [4](#)
20. Li, Z., Niklaus, S., Snavely, N., Wang, O.: Neural scene flow fields for space-time view synthesis of dynamic scenes. In: CVPR. pp. 6498–6508 (2021) [1](#), [2](#), [3](#), [5](#), [9](#), [10](#), [11](#), [12](#), [13](#)
21. Li, Z., Wang, Q., Cole, F., Tucker, R., Snavely, N.: Dynibar: Neural dynamic image-based rendering. In: CVPR. pp. 4273–4284 (2023) [1](#), [2](#), [5](#), [10](#)
22. Liu, Y.L., Gao, C., Meuleman, A., Tseng, H.Y., Saraf, A., Kim, C., Chuang, Y.Y., Kopf, J., Huang, J.B.: Robust dynamic radiance fields. In: CVPR. pp. 13–23 (2023) [5](#), [9](#), [11](#), [12](#)
23. Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehmman, A., Sheikh, Y.: Neural volumes: Learning dynamic renderable volumes from images. *ACM Transactions on Graphics* **38**(4CD), 65.1–65.14 (2019) [4](#)
24. Luo, X., Peng, J., Xian, K., Wu, Z., Cao, Z.: Defocus to focus: Photo-realistic bokeh rendering by fusing defocus and radiance priors. *Information Fusion* **89**, 320–335 (2023) [4](#)
25. Ma, L., Li, X., Liao, J., Zhang, Q., Wang, X., Wang, J., Sander, P.V.: Deblurrer: Neural radiance fields from blurry images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12861–12870 (2022) [2](#), [3](#), [4](#), [8](#), [11](#)
26. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* **38**(4), 1–14 (2019) [4](#)
27. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 405–421. Springer (2020) [2](#), [4](#)
28. Pan, L., Chowdhury, S., Hartley, R., Liu, M., Zhang, H., Li, H.: Dual pixel exploration: Simultaneous depth estimation and image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4340–4349 (2021) [4](#)
29. Park, J., Tai, Y.W., Cho, D., So Kweon, I.: A unified approach of multi-scale deep and hand-crafted features for defocus estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1736–1745 (2017) [4](#)
30. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: ICCV. pp. 5865–5874 (2021) [5](#)
31. Park, K., Sinha, U., Hedman, P., Barron, J.T., Bouaziz, S., Goldman, D.B., Martin-Brualla, R., Seitz, S.M.: Hypernerf: a higher-dimensional representation for topologically varying neural radiance fields. *ACM TOG* **40**(6), 1–12 (2021) [2](#), [5](#), [11](#), [12](#)
32. Peng, J., Cao, Z., Luo, X., Lu, H., Xian, K., Zhang, J.: Bokehme: When neural rendering meets classical rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16283–16292 (2022) [4](#), [6](#), [11](#)

33. Peng, J., Zhang, J., Luo, X., Lu, H., Xian, K., Cao, Z.: Mpib: An mpi-based bokeh rendering framework for realistic partial occlusion effects. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VI. pp. 590–607. Springer (2022) [4](#), [6](#)
34. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: CVPR. pp. 10318–10327 (2021) [1](#), [5](#)
35. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12179–12188 (2021) [10](#)
36. Ruan, L., Chen, B., Li, J., Lam, M.: Learning to deblur using light field generated and real defocus images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16304–16313 (2022) [4](#)
37. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4104–4113 (2016) [11](#)
38. Sheng, Y., Yu, Z., Ling, L., Cao, Z., Zhang, C., Lu, X., Xian, K., Lin, H., Benes, B.: Dr. bokeh: Differentiable occlusion-aware bokeh rendering. arXiv preprint arXiv:2308.08843 (2023) [4](#)
39. Shi, J., Xu, L., Jia, J.: Just noticeable defocus blur detection and estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 657–665 (2015) [4](#)
40. Sitzmann, V., Thies, J., Heide, F., Nießner, M., Wetzstein, G., Zollhofer, M.: Deepvoxels: Learning persistent 3d feature embeddings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2437–2446 (2019) [4](#)
41. Sitzmann, V., Zollhofer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 32. Curran Associates, Inc. (2019), <https://proceedings.neurips.cc/paper/2019/file/b5dc4e5d9b495d0196f61d45b26ef33e-Paper.pdf> [4](#)
42. Son, H., Lee, J., Cho, S., Lee, S.: Single image defocus deblurring using kernel-sharing parallel atrous convolutions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2642–2650 (2021) [4](#), [11](#), [12](#)
43. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16. pp. 402–419. Springer (2020) [10](#)
44. Tretschk, E., Tewari, A., Golyanik, V., Zollhöfer, M., Lassner, C., Theobalt, C.: Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In: ICCV. pp. 12959–12970 (2021) [5](#)
45. Wadhwa, N., Garg, R., Jacobs, D.E., Feldman, B.E., Kanazawa, N., Carroll, R., Movshovitz-Attias, Y., Barron, J.T., Pritch, Y., Levoy, M.: Synthetic depth-of-field with a single-camera mobile phone. ACM Transactions on Graphics (TOG) **37**(4), 1–13 (2018) [4](#), [6](#)
46. Wang, C., Eckart, B., Lucey, S., Gallo, O.: Neural trajectory fields for dynamic novel view synthesis. arXiv preprint arXiv:2105.05994 (2021) [5](#)
47. Wang, C., MacDonald, L.E., Jeni, L.A., Lucey, S.: Flow supervision for deformable nerf. In: CVPR. pp. 21128–21137 (2023) [5](#)
48. Wang, L., Shen, X., Zhang, J., Wang, O., Lin, Z., Hsieh, C.Y., Kong, S., Lu, H.: Deeplens: Shallow depth of field from a single image. ACM Transactions on Graphics (TOG) **37**(6), 1–11 (2018) [4](#)

49. Wang, Y., Shi, M., Li, J., Huang, Z., Cao, Z., Zhang, J., Xian, K., Lin, G.: Neural video depth stabilizer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9466–9476 (2023) [3](#), [11](#)
50. 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. pp. 1718–1729 (2022) [3](#), [4](#)
51. Xian, W., Huang, J.B., Kopf, J., Kim, C.: Space-time neural irradiance fields for free-viewpoint video. In: CVPR. pp. 9421–9431 (2021) [5](#)
52. Xiao, L., Kaplanyan, A., Fix, A., Chapman, M., Lanman, D.: Deepfocus: Learned image synthesis for computational displays. *ACM Transactions on Graphics (TOG)* **37**(6), 1–13 (2018) [4](#)
53. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5728–5739 (2022) [4](#)
54. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 586–595 (2018) [12](#)
55. 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**, 1–16 (2019) [4](#)