

DyBluRF: Dynamic Neural Radiance Fields from Blurry Monocular Video

Huiqiang Sun¹ Xingyi Li¹ Liao Shen¹ Xinyi Ye¹ Ke Xian² Zhiguo Cao^{1*}

¹School of AIA, Huazhong University of Science and Technology

²School of EIC, Huazhong University of Science and Technology

{shq1031, xingyi_li, leoshen, xinyiye, kxian, zgcao}@hust.edu.cn

<https://huiqiang-sun.github.io/dyblurf>

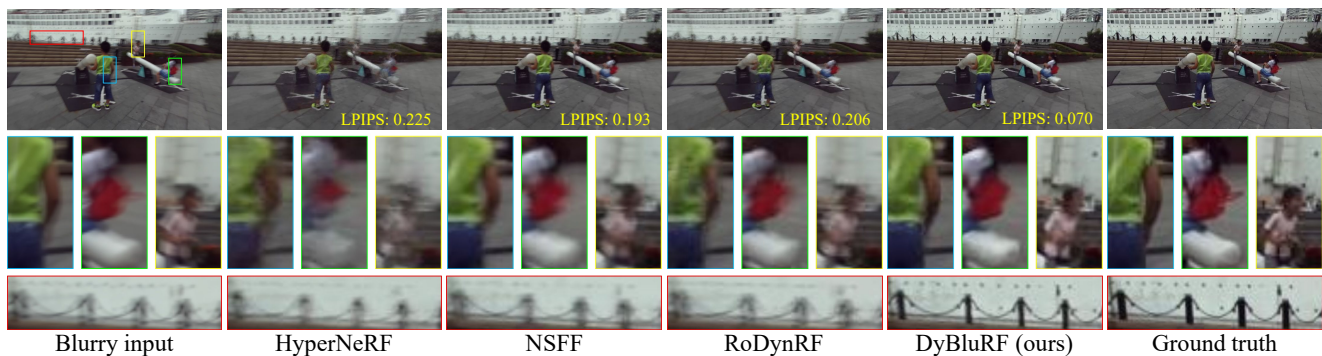


Figure 1. Given a monocular video capturing a dynamic scene with motion blur, our proposed method, DyBluRF, effectively synthesizes high-quality and sharp novel views compared to previous dynamic NeRF approaches that often yield low-quality and blurry results.

Abstract

Recent advancements in dynamic neural radiance field methods have yielded remarkable outcomes. However, these approaches rely on the assumption of sharp input images. When faced with motion blur, existing dynamic NeRF methods often struggle to generate high-quality novel views. In this paper, we propose DyBluRF, a dynamic radiance field approach that synthesizes sharp novel views from a monocular video affected by motion blur. To account for motion blur in input images, we simultaneously capture the camera trajectory and object Discrete Cosine Transform (DCT) trajectories within the scene. Additionally, we employ a global cross-time rendering approach to ensure consistent temporal coherence across the entire scene. We curate a dataset comprising diverse dynamic scenes that are specifically tailored for our task. Experimental results on our dataset demonstrate that our method outperforms existing approaches in generating sharp novel views from motion-blurred inputs while maintaining spatial-temporal consistency of the scene.

*Corresponding author.

1. Introduction

The research on novel view synthesis for dynamic scenes has gained substantial importance in the community, with practical applications in a wide range of AR/VR contexts. Recent breakthroughs can be attributed to Neural Radiance Fields (NeRFs) [27]. NeRFs have made a profound impact on novel view synthesis, and recent years have witnessed rapid advancements in NeRF, including several works dedicated to dynamic scene representation [10, 21, 24, 32, 34]. Despite producing high-quality novel views, these dynamic NeRF methods may exhibit significant performance degradation when presented with inputs containing motion blur, as shown in Fig. 1. Recently, some NeRF-based methods have emerged to tackle the motion deblurring of static scenes [18, 19, 25, 46]. However, these methods cannot address blurring caused by object motion and lack the capability to effectively represent dynamic scenes.

To the best of our knowledge, there is no NeRF-based method dedicated to dynamic scene deblurring. Yet, in real-life video captures of dynamic scenarios, extended camera exposure time often introduces motion blur. Consequently, obtaining a sharp dynamic radiance field from inputs affected by motion blur remains a pressing problem. Existing dynamic NeRF methods heavily rely on sharp inputs

to accurately represent dynamic scenes. When motion blur is present, the modulation of inter-frame relationships becomes unreliable, leading to difficulties in representing object motion and maintaining temporal consistency in dynamic scenes. Moreover, the presence of blur in input data degrades certain data priors like depth and optical flow, leading to inaccurate constraining of scene geometry. These challenges significantly hinder the construction of sharp dynamic NeRF representations from blurry inputs.

In this work, we propose DyBluRF, a model specifically designed to take monocular videos of dynamic scenes with motion blur as input and generate sharp novel views. Motion blur in dynamic scenes typically results from both camera and object movements. To model camera motion, drawing inspiration from BAD-NeRF [46], we evenly discretize the exposure time into multiple timestamps and learn camera poses at these timestamps. For object motion representation, establishing inter-frame relationships, such as directly predicting 3D scene flow, is unreliable under conditions of blurry input, thus failing to accurately present dynamic scenes. Instead, we employ an MLP to predict the global DCT trajectories of scene objects [43], thereby simulating object motion blur by expressing the DCT trajectories within an exposure time.

To enhance the temporal consistency representation in dynamic scenes, we adopt a cross-time rendering approach based on the input blurry images. This method uses the predicted DCT trajectory to model scene correlation across multiple input views and employs deformed rays for rendering by integrating scene information from other frames into the target frame. To ensure a globally consistent temporal expression, cross-time rendering is not limited to adjacent frames but instead performed across the entire temporal range. At the initial stages of training, frames close to the target frame are selected for cross-time rendering. As the training progresses, frames further away from the target frame are included until encompass the entire temporal range. This cross-time rendering approach further strengthens the representation of temporal consistency.

Furthermore, it is not reasonable to directly use depth and optical flow maps generated from blurry images to constrain the model, as the predictions of depth and optical flow can become inaccurate due to blurriness. Therefore, we introduce data-driven priors to mitigate the impact of inaccuracies in depth and optical flow from blurry images.

Given the absence of existing dynamic scene blurring datasets based on NeRF, we curate a collection of dynamic scenes sourced from the Stereo Blur Dataset [57]. We use these scenes for comparative analysis of our approach against existing dynamic NeRF methods, demonstrating its effective handling of input images with motion blur and generation of high-quality novel views. We also preprocess the input blurry images using state-of-the-art single-image

deblurring and dynamic scene video deblurring methods before feeding them into existing dynamic NeRF methods. Experimental results demonstrate that our method still outperforms existing dynamic NeRF methods. Additionally, we conduct ablation studies to validate the effectiveness of each component of our approach. In summary, our key contributions include:

- We introduce DyBluRF, the first dynamic NeRF model specifically designed to effectively tackle motion-blurred monocular video.
- We propose a dynamic scene representation method based on DCT, simulating the motion blur induced by object movement along the DCT trajectory.
- We present a cross-time rendering method that ensures temporal consistency by modeling scene relationships across different exposure times.
- We propose an effective approach that incorporates depth maps and optical flow derived from blurry images to impose constraints on NeRF.

2. Related Work

Image Deblurring. The main goal of image deblurring is to restore a sharp image from a blurry one. From the perspective of input data types, current motion deblurring methods can be categorized into single-image deblurring and multi-image/video deblurring. In the case of single-image deblurring, Some traditional methods employ unified probabilistic model [37], normalized gradient sparsity [15], or model rotational velocity of the camera during exposure time [48] to mitigate motion blur. Recently, numerous deep learning-based image deblurring methods have emerged [16, 17, 29, 40, 52], accompanied by the introduction of several datasets dedicated to image deblurring [38]. These methods often achieve superior results compared to traditional techniques.

In contrast to single-image deblurring, video deblurring requires establishing temporal consistency among different frames to ensure temporal smoothness. Some approaches employ optical flow information to establish inter-frame relationships [12, 30]. However, the optical flow in blurry videos often lacks accuracy. With the advance of deep learning, many techniques based on CNNs [39], RNNs [55], or attention mechanisms [53] have been proposed to address the problem of video deblurring. While these methods yield excellent results, they predominantly operate in image space and cannot capture scene geometry information in 3D space. Our approach leverages NeRF to represent 3D information of dynamic scenes implicitly and establishes spatial-temporal consistency to achieve deblurring.

Neural Radiance Field. NeRF [27] has demonstrated impressive results in novel view synthesis, and subsequent works have aimed to enhance it from various perspectives, including fundamental enhancements [2, 3], fast training

and rendering [6, 7, 9, 28], pose estimation [4, 23, 47], or sparse viewpoint input [50, 51]. Since vanilla NeRF imposes high requirements on input image quality, several follow-up methods have explored generating high-quality novel views in degraded input conditions, such as in-the-wild images [26], input images with burst noise [33], or low dynamic range (LDR) inputs with varying exposures [11]. Recently, several NeRF-based approaches have been devised to handle blurry inputs. Deblur-NeRF [25] and DP-NeRF [18] restore sharp radiance fields by modeling a blur kernel from various blur types. BAD-NeRF [46], which is most relevant to our work, models motion blur by estimating camera trajectory. ExBluRF [19] leverages a voxel-based radiance field to handle extreme motion blur. While these methods can remove blur from input images and generate high-quality novel views, they cannot handle dynamic scenes. Our approach can effectively reduce motion blur caused by both camera and object motion, generating novel views of dynamic scenes.

Dynamic Neural Radiance Field. Many recent efforts have extended NeRF to represent dynamic scenes. These methods input multi-view videos [20] or a monocular video [22, 49] of dynamic scenes to achieve space-time novel view synthesis. To ensure temporal consistency, some approaches represent dynamic scenes by deforming the canonical space points [31, 32, 34, 42, 45]. Other methods leverage optical flow constraints to establish scene coherence between adjacent frames [10, 21, 24, 44]. In particular, NSFF [21] employs neural scene flow fields to represent complex object motions in dynamic scenes; RoDynRF [24] jointly learns camera parameters to enhance model robustness. Additionally, many methods are currently dedicated to fast training and rendering of dynamic NeRF [1, 5, 8]. However, these dynamic NeRF methods require sharp input images and often struggle to synthesize reliable novel views in the presence of motion blur in the input images.

3. Method

We propose DyBluRF, a dynamic radiance field approach synthesizing sharp novel views from a monocular video affected by motion blur. The overall pipeline is illustrated in Fig. 2. In this section, we first provide background information on neural radiance fields and volume rendering in Section 3.1. We then elaborate on modeling genuine physical image formation process of motion blur in Section 3.2. Next, Section 3.3 outlines our approach for representing dynamic scenes and ensuring temporal consistency within the scene. Finally, we provide details of the training and optimization process in Section 3.4.

3.1. Preliminary: Neural Radiance Field

We begin by describing Neural Radiance Fields (NeRFs) that our method builds upon. NeRF represents a scene as an

implicit function that inputs a spatial point location $\mathbf{x} \in \mathbb{R}^3$ and viewing direction $\mathbf{d} \in \mathbb{R}^3$, and outputs corresponding volume density σ and color \mathbf{c} :

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

To enable learning scene details, point location \mathbf{x} and viewing direction \mathbf{d} are transformed into a high-dimensional space through positional embedding $\gamma(\cdot)$ before sending to NeRF. We also employ this trick in our methods but omit to mention the position encoding in our writing for clarity.

NeRF adopts classical volume rendering [13] to render RGB images from any given camera pose. For a ray \mathbf{r} emitted from camera center through a given pixel on the image plane, its color is computed as:

$$\mathbf{C}(\mathbf{r}) = \int_{s_n}^{s_f} T(s) \sigma(\mathbf{r}(s)) \mathbf{c}(\mathbf{r}(s), \mathbf{d}) ds, \quad (2)$$

where

$$T(s) = \exp\left(-\int_{s_n}^s \sigma(\mathbf{r}(m)) dm\right) \quad (3)$$

is accumulated transmittance along ray \mathbf{r} which denotes the probability that the ray will not hit any object from s_n to s .

3.2. Motion Blur Formulation

The physical process of motion blur can be considered as the result of camera or object motion within exposure time τ . The mathematical modeling of the motion blur generation process can be written as:

$$\mathbf{B}(\mathbf{x}) = g \int_0^{\tau} \mathbf{C}_t(\mathbf{x}) dt, \quad (4)$$

where $\mathbf{C}_t(\mathbf{x})$ denotes the sharp image captured at time t , $\mathbf{B}(\mathbf{x})$ is the blurry image, g is a normalization factor. To approximate the integral operation, we discretize the exposure time into n timestamps uniformly and consider the captured blurry image as the average of n sharp images during the exposure time:

$$\mathbf{B}(\mathbf{x}) \approx \frac{1}{n} \sum_{l=1}^n \mathbf{C}_l(\mathbf{x}). \quad (5)$$

In theory, as n increases, the simulated blur generation process becomes more close to a genuine one.

According to Eq. 5, generating a blurry image demands obtaining sharp images at each timestamp within the exposure time. Therefore, it becomes crucial to model both camera and object motion during the exposure time. For the object motion representation, we will provide a detailed explanation in Section 3.3. For the camera motion, we adopt a strategy that jointly learns camera parameters and NeRF. Specifically, given N input blurry images, we initialize the

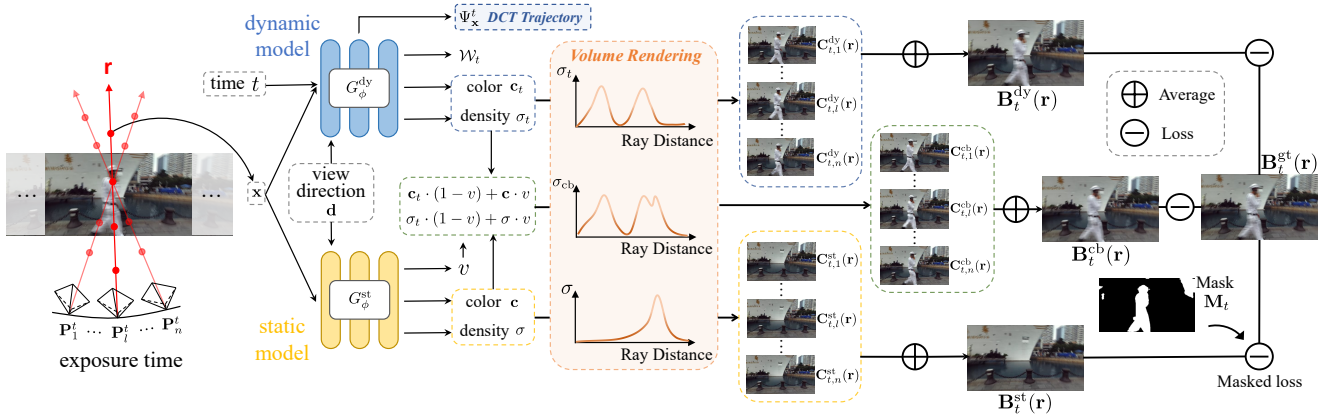


Figure 2. **Overall pipeline.** To model the motion blur of input images, we discretize the exposure time into n timestamps. Subsequently, we perform ray sampling for the same pixel based on the camera poses for each timestamp. We employ two MLPs to represent dynamic scenes. The dynamic model takes spatial point coordinates \mathbf{x} , viewing direction \mathbf{d} , and time t as inputs and predicts color \mathbf{c}_t , volume density σ_t , DCT coefficients $\Psi_{\mathbf{x}}^t$, and disocclusion weights \mathcal{W}_t . The static model only takes \mathbf{x} and \mathbf{d} as inputs and predicts color \mathbf{c} , volume density σ , and a blending weight v for blending static and dynamic results. After obtaining colors and volume densities for static, dynamic, and blended scenes, we use volume rendering to compute pixel RGB values $\mathbf{C}(\mathbf{r})$. We average the RGB values for the same pixel within the exposure time to obtain the predicted blurry image, and calculate losses with the input blurry image. For dynamic and blended images, losses are directly computed against the input image. For the static loss, we use a mask image to calculate the static regions only.

camera poses \mathbf{P}_1^i and \mathbf{P}_n^i at the beginning and end timestamps of exposure time i , where $i \in \{1, \dots, N\}$. Leveraging the short exposure time, we utilize linear interpolation to obtain camera poses for intermediate timestamps. Following the ideas of BAD-NeRF [46], we employ Lie-algebra of $SE(3)$ to interpolate the intermediate camera poses \mathbf{P}_l^i , $l \in \{2, \dots, n-1\}$.

3.3. Temporal Consistency Modeling

Dynamic radiance field. Similar to existing dynamic NeRF methods, we employ an MLP to represent dynamic scenes. This MLP takes the input spatial point location \mathbf{x} , viewing direction \mathbf{d} , and time t , and outputs the corresponding color and volume density: $G_{\phi}^{\text{dy}} : (\mathbf{c}_t, \sigma_t) = G_{\phi}^{\text{dy}}(\mathbf{x}, \mathbf{d}, t)$, where G_{ϕ}^{dy} indicates the MLP parameterized by ϕ . Handling a monocular video where scene information at different times is limited to individual frames poses an ill-posed problem. Establishing temporal consistency in dynamic scenes becomes essential to resolve this issue. Although employing scene flow information between adjacent frames is a direct approach [21, 24], accurately determining scene motion within the exposure time using adjacent blurry images remains challenging. To enforce a more reliable temporal consistency, we establish relationships across global time intervals rather than relying solely on adjacent frames. Leveraging the efficiency of the Discrete Cosine Transform (DCT) in representing complex motion trajectories in a smooth and concise manner [43, 44], we choose to model the object DCT trajectory to express dynamic scenes.

Given a sequence of N blurry images, the exposure time

for each image can be divided into n timestamps, resulting in a total of $n \times N$ timestamps. The trajectory with DCT coefficients $\Psi \in \mathbb{R}^{3K}$ can be expressed as:

$$\mathcal{T}(t) = \sqrt{\frac{2}{nN}} \sum_{k=1}^K \Psi(k) \cos\left(\frac{\pi}{2nN}(2t+1)k\right), \quad (6)$$

where $\mathcal{T} \in \mathbb{R}^3$ represents the time-dependent trajectory, $\Psi(k) \in \mathbb{R}^3$ denotes the k -th coefficient in 3D space, K is a hyperparameter. According to Eq. 6, for a 3D point \mathbf{x} at time t , its global time motion trajectory $\mathcal{T}_{\mathbf{x}}^t$ can be modeled using the corresponding DCT coefficient $\Psi_{\mathbf{x}}^t$, and $\mathcal{T}_{\mathbf{x}}^t(t')$ represents its 3D coordinates at time t' . We use the MLP to predict the DCT coefficients. Thus, the dynamic radiance field can be rewritten as $G_{\phi}^{\text{dy}} : (\mathbf{c}_t, \sigma_t, \Psi_{\mathbf{x}}^t) = G_{\phi}^{\text{dy}}(\mathbf{x}, \mathbf{d}, t)$.

To handle occlusion in dynamic scenes, we adopt a similar approach to [21], where the MLP additionally outputs a disocclusion weight $\mathcal{W}_t = (w_{\text{fw}}, w_{\text{bw}})$. However, unlike [21], our method predicts disocclusion weights to represent the occlusion possibility to the corresponding timestamps within the subsequent or preceding exposure time. In summary, our dynamic radiance fields G_{ϕ}^{dy} can be expressed as:

$$(\mathbf{c}_t, \sigma_t, \Psi_{\mathbf{x}}^t, \mathcal{W}_t) = G_{\phi}^{\text{dy}}(\mathbf{x}, \mathbf{d}, t). \quad (7)$$

To supervise the model using input images, we need to obtain the blurry image at each frame through the model. For the i -th blurry image $\mathbf{B}_i(\mathbf{r})$, we calculate the timestamps t_l^i , $l \in \{1, \dots, n\}$ within the i -th camera exposure time. As per Section 3.2, the camera pose for each timestamp is denoted as \mathbf{P}_l^i . Using the MLP depicted in Eq. 7

and the volume rendering from Eq. 2, we derive predicted sharp images $\mathbf{C}_{i,l}^{\text{dy}}(\mathbf{r})$ of timestamps l within the i -th exposure time. Then, we average these sharp images to obtain the predicted blurry image $\mathbf{B}_i^{\text{dy}}(\mathbf{r})$:

$$\mathbf{B}_i^{\text{dy}}(\mathbf{r}) = \frac{1}{n} \sum_{l=1}^n \mathbf{C}_{i,l}^{\text{dy}}(\mathbf{r}), \quad i \in \{1, \dots, N\}, \quad (8)$$

and calculate rendering loss with input blurry image $\mathbf{B}_i^{\text{gt}}(\mathbf{r})$:

$$\mathcal{L}_{\text{RGB}}^{\text{dy}} = \|\mathbf{B}_i^{\text{dy}}(\mathbf{r}) - \mathbf{B}_i^{\text{gt}}(\mathbf{r})\|_2^2. \quad (9)$$

Static model integration. Similar to many previous dynamic NeRF methods [21, 24], we also employ an additional MLP to enhance the representation of static regions:

$$(\mathbf{c}, \sigma, v) = G_{\phi}^{\text{st}}(\mathbf{x}, \mathbf{d}), \quad (10)$$

where v denotes the blending weight used to combine the results from static and dynamic models:

$$\begin{aligned} \mathbf{c}_{\text{combine}} &= \mathbf{c}_t \cdot (1 - v) + \mathbf{c} \cdot v, \\ \sigma_{\text{combine}} &= \sigma_t \cdot (1 - v) + \sigma \cdot v. \end{aligned} \quad (11)$$

We utilize Eq. 2 to obtain static sharp images $\mathbf{C}_{i,l}^{\text{st}}(\mathbf{r})$ and combined sharp images $\mathbf{C}_{i,l}^{\text{cb}}(\mathbf{r})$, and we average these images in an exposure time to obtain the predicted static blurry image $\mathbf{B}_i^{\text{st}}(\mathbf{r})$ and combined blurry image $\mathbf{B}_i^{\text{cb}}(\mathbf{r})$. Subsequently, we apply rendering loss from both static and combined results to constrain the MLPs:

$$\begin{aligned} \mathcal{L}_{\text{RGB}}^{\text{cb}} &= \|\mathbf{B}_i^{\text{cb}}(\mathbf{r}) - \mathbf{B}_i^{\text{gt}}(\mathbf{r})\|_2^2, \\ \mathcal{L}_{\text{RGB}}^{\text{st}} &= (1 - \mathbf{M}_i) \cdot \|\mathbf{B}_i^{\text{st}}(\mathbf{r}) - \mathbf{B}_i^{\text{gt}}(\mathbf{r})\|_2^2. \end{aligned} \quad (12)$$

where \mathbf{M}_i denotes the motion mask of input frame i . The inclusion of static model enhances the ability to represent static scenes.

Cross-time rendering. If we only rely on Eq. 9 for reconstruction, the model may inadequately capture the temporal consistency in dynamic scenes, potentially leading to overfitting to input images and failing to produce high-quality novel views. To mitigate this issue, we utilize DCT trajectories to establish consistency between the i -th frame and other frames, as shown in Fig. 3. Establishing temporal consistency between neighboring frames plays a critical role in the dynamic radiance field due to their strong correlation [10, 22]. However, blurry images relate to multiple images within the exposure time rather than a single image. Therefore, we express temporal consistency by establishing relationships between corresponding timestamps across different exposure times. Specifically, for a timestamp t_l^i in an exposure time from frame i , the scene flow of a spatial point \mathbf{x} from t_l^i to timestamp t_l^j can be represented as:

$$f_{\mathbf{x}}^{i,l}(j) = \mathcal{T}_{\mathbf{x}}^{t_l^j} - \mathcal{T}_{\mathbf{x}}^{t_l^i}, \quad (13)$$

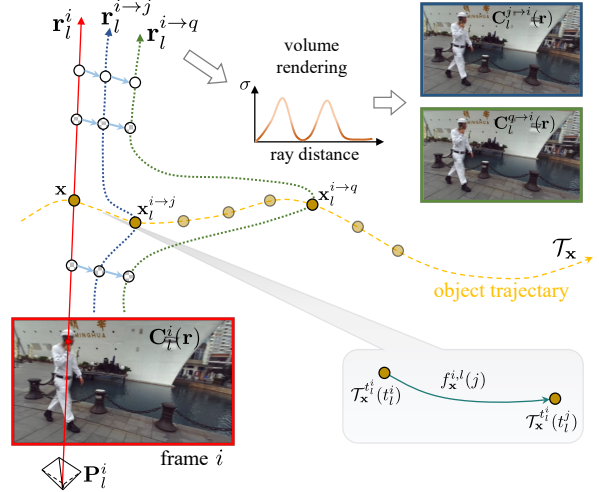


Figure 3. **Cross-time rendering.** We use cross-time rendering to ensure temporal consistency in dynamic scene representation. We render the target timestamp image $\mathbf{C}_l^i(\mathbf{r})$ in frame i utilizing scene information predicted from the corresponding timestamp in other frames. In addition to selecting adjacent frames $j \in \mathcal{N}(i)$, we also choose an extra frame q from the global time to ensure the global consistency of the DCT trajectory.

where $j \in \mathcal{N}(i)$ denotes the neighbor frames of i . The corresponding 3D point of \mathbf{x} at t_l^j can be computed as $\mathbf{x}_l^{i \rightarrow j} = \mathbf{x} + f_{\mathbf{x}}^{i,l}(j)$. Establishing such relationships for every point on a ray \mathbf{r} can obtain a warped ray $\mathbf{r}_l^{i \rightarrow j}$ from t_l^i to t_l^j . Subsequently, we utilize volume rendering to obtain the color value of the warped ray $\mathbf{C}_l^{j \rightarrow i}(\mathbf{r})$:

$$\mathbf{C}_l^{j \rightarrow i}(\mathbf{r}) = \int_{s_n}^{s_f} T_{t_l^j}(s) \sigma_{t_l^j}(\mathbf{r}_l^{i \rightarrow j}(s)) \mathbf{c}_{t_l^j}(\mathbf{r}_l^{i \rightarrow j}(s), \mathbf{d}_{t_l^j}) ds. \quad (14)$$

Then we average the images within the exposure time to yield the blurred image $\mathbf{B}^{j \rightarrow i}(\mathbf{r})$, and we calculate the cross-time rendering loss between $\mathbf{B}^{j \rightarrow i}(\mathbf{r})$ and $\mathbf{B}_i^{\text{gt}}(\mathbf{r})$:

$$\mathbf{B}^{j \rightarrow i}(\mathbf{r}) = \frac{1}{n} \sum_{l=1}^n \mathbf{C}_l^{j \rightarrow i}(\mathbf{r}), \quad (15)$$

$$\mathcal{L}_{\text{cross}} = \sum_{j \in \mathcal{N}(i)} \mathbf{W}^{j \rightarrow i}(\mathbf{r}) \|\mathbf{B}^{j \rightarrow i}(\mathbf{r}) - \mathbf{B}_i^{\text{gt}}(\mathbf{r})\|_2^2, \quad (16)$$

where $\mathbf{W}^{j \rightarrow i}(\mathbf{r})$ represents a motion disocclusion weight, which is the average outcome of n disocclusion weights $\mathbf{W}_l^{j \rightarrow i}(\mathbf{r})$ computed by \mathcal{W}_t in Eq. 7. For the specifics of disocclusion weight settings, we offer detailed information in the supplementary material. Through the aforementioned cross-time rendering, the DCT coefficients Ψ predicted for the same point at different times remain consistent, which preserves the temporal consistency of the scene.

In addition to using adjacent frames $j \in \mathcal{N}(i)$ for cross-time rendering, we include an additional frame q for ensuring global time consistency. During the early training stages, q is selected to be close to i . As training progresses, q is gradually chosen further away from i until it encompasses the entire global time range. We employ a similar rendering process to obtain the blurry image $\mathbf{B}^{q \rightarrow i}(\mathbf{r})$ and compute cross-time rendering loss with input image. In summary, the cross-time rendering loss can be written as:

$$\mathcal{L}_{\text{cross}} = \sum_{j \in \{\mathcal{N}(i), q\}} \mathbf{W}^{j \rightarrow i}(\mathbf{r}) \|\mathbf{B}^{j \rightarrow i}(\mathbf{r}) - \mathbf{B}_i^{\text{gt}}(\mathbf{r})\|_2^2. \quad (17)$$

3.4. Training & Optimization

Data-driven priors. We introduce data-driven priors $\mathcal{L}_{\text{data}}$ to constrain the scene geometry. We utilize RAFT [41] and MiDaS [35] to obtain ground truth optical flow and depth for the input blurry images. Intuitively, we only need to calculate the depth maps for each timestamp, and optical flow maps from each timestamp to the next exposure time, subsequently averaging them within a single exposure time and constraining them with the ground truth. However, we encountered issues where the depth and optical flow do not have a translucent effect in the blurry edges, as opposed to observed in RGB data (we show a depth map as an example in Fig. 4(a)). Therefore, using simple averaging is unsuitable for constraining depth and optical flow. To address this issue, we implement Extreme Value Constraints (EVC) for both depth and optical flow maps. Specifically, after calculating the depth maps $\{\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_n\}$ and optical flow maps $\{\mathbf{F}_1, \mathbf{F}_2, \dots, \mathbf{F}_n\}$ projected from 3D scene flow within an exposure time, we derive the predicted blurry result by taking the minimum depth value and maximum optical flow value, and calculate the loss with ground truth:

$$\begin{aligned} \mathbf{D}(\mathbf{r}) &= \min\{\mathbf{D}_1(\mathbf{r}), \mathbf{D}_2(\mathbf{r}), \dots, \mathbf{D}_n(\mathbf{r})\}, \\ \mathbf{F}(\mathbf{r}) &= \max\{\mathbf{F}_1(\mathbf{r}), \mathbf{F}_2(\mathbf{r}), \dots, \mathbf{F}_n(\mathbf{r})\}, \end{aligned} \quad (18)$$

$$\mathcal{L}_{\text{data}} = \|\mathbf{D}(\mathbf{r}) - \mathbf{D}_{\text{gt}}(\mathbf{r})\|_1 + \|\mathbf{F}(\mathbf{r}) - \mathbf{F}_{\text{gt}}(\mathbf{r})\|_1. \quad (19)$$

After applying EVC for data-driven constraints, a sharper result can be generated for each timestamp (as shown in Fig. 4(b)). Please refer to supplementary material for principles and technical details of EVC.

Scene flow modeling. For a more accurate representation of scene motion, we use the predicted DCT coefficients Ψ to calculate the scene flow and impose constraints using \mathcal{L}_{sf} . Similar to NSFF [21], \mathcal{L}_{sf} consists of three components: cycle consistency of the scene flow, spatial-temporal smoothness constraints, and L1 regularization to the scene flow.

Learning DCT basis. We model the motion trajectories using Eq. 6; however, we discovered that employing a learnable DCT basis produced superior outcomes. To achieve

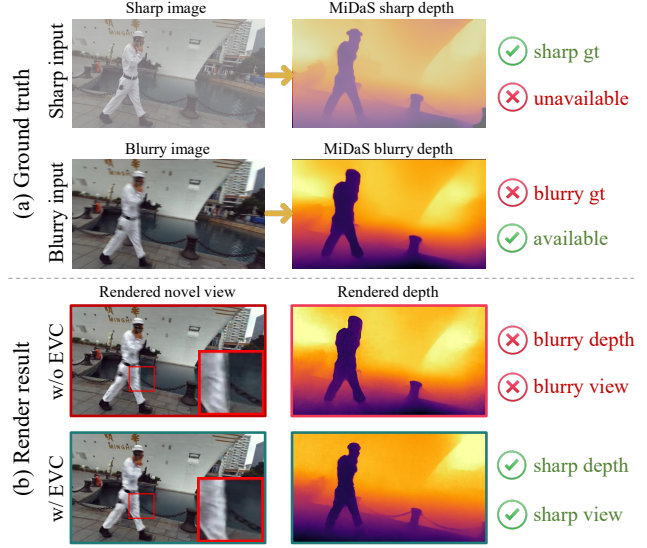


Figure 4. (a) We utilize depth map from MiDaS [35] as the supervising signal for optimizing the model. Compared with sharp inputs, blurry images are available but lead to inaccurate depth ground truth. In the example, MiDaS interprets the blurry edges of the person as foreground, causing the person on the depth map to appear ‘fatter’ than in the ideal sharp one. (b) Using the MiDaS blurry depth to optimize the model directly may predict inaccurate depth maps and distorted novel views. To mitigate this issue, we employ EVC to enhance our model for predicting sharp depth maps and novel views from blurry images.

this, we initialize the DCT basis using the expression $\cos\left(\frac{\pi}{2 * n * N} (2t + 1)k\right)$, and treat them as learnable embeddings, facilitating joint learning with the models throughout the training process.

Implementation details. In experiments, we set $n = 7$ and $K = 6$. We employ the Adam optimizer [14] to jointly optimize the static and dynamic MLPs, camera parameters, and DCT basis. The learning rates are set to 5×10^{-4} for MLPs, 1×10^{-3} for camera parameters, and 1.25×10^{-4} for DCT basis. We train each scene for 300k iterations, which takes approximately two days on a single NVIDIA RTX A6000 GPU.

4. Experiments

4.1. Datasets

Since there is a lack of datasets specifically designed for addressing motion blur in dynamic neural radiance fields, we have curated a collection of dynamic scenes with motion blur from an existing Stereo Blur Dataset [57]. This dataset was captured using a ZED stereo camera, containing sharp stereo image sequences along with their corresponding blurry image counterparts. Details about the generation of blur and other dataset information can be found



Figure 5. **Qualitative comparisons against all baselines.** Compared to existing dynamic NeRF methods, our method generates novel view images that are more faithful to the ground truth images, with less blur in both static and dynamic regions.

Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
HyperNeRF [32]	20.15	0.744	0.221	[52] + [21]	23.17	0.839	0.143
DVS [10]	22.30	0.811	0.193	[56] + [21]	<u>23.79</u>	<u>0.854</u>	<u>0.125</u>
NSFF [21]	<u>23.33</u>	<u>0.834</u>	<u>0.181</u>	[52] + [24]	20.76	0.787	0.162
RoDynRF [24]	19.28	0.748	0.217	[56] + [24]	20.88	0.794	0.151
DyBluRF (ours)	25.66	0.895	0.079	DyBluRF	25.66	0.895	0.079

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

in [57], and we also provide additional explanations in supplementary materials. To obtain accurate camera parameters for training neural radiance fields, we require image sequences of dynamic scenes with sufficient motion parallax. We collected 6 dynamic scenes suitable for our task from this dataset. These scenes exhibit significant motion blur caused by both camera and object movements. In our experiments, we employ COLMAP [36] to acquire the camera parameters for the input images. For each scene, we extract 24 frames from the original video, utilizing the left blurry image sequences for training and the corresponding right sharp image sequences for testing.

4.2. Baselines and Error Metrics

We compare our method against several state-of-the-art monocular dynamic NeRF methods. Specifically, the comparison involves a canonical space-based method, HyperNeRF [32], two scene flow-based methods, NSFF [21] and DVS [10], and a robust dynamic NeRF method, RoDynRF [24]. Furthermore, to demonstrate the superiority of our approach over 2D image deblurring, we select the current state-of-the-art single-image deblurring method

Restormer [52] and a dynamic scene video deblurring method JCD [56] to perform deblurring preprocessing on the input blurry images. We then train NSFF and RoDynRF on the preprocessed images for comparison.

Similar to most existing dynamic NeRF methods, we use PSNR, SSIM, and LPIPS [54] to quantitatively evaluate the quality of novel views generated by different approaches.

4.3. Comparisons

We conduct a quantitative evaluation on the scenes collected from the Stereo Blur Dataset. The average results across all scenes are shown in Tab. 1 (left), clearly demonstrating the superior performance of our method compared to the existing dynamic NeRF methods. We also provide qualitative comparisons of novel view results generated by different methods, as shown in Fig. 5. Notably, existing dynamic NeRF methods cannot handle motion blur inputs, resulting in low-quality novel views in both dynamic and static regions. In contrast, our method can effectively manage inputs with motion blur and produce sharp and high-quality novel views. Furthermore, to demonstrate the superiority of our method in modeling motion blur in 3D space, we first preprocess the input blurry images with 2D deblurring and then train dynamic NeRF models, e.g., RoDynRF, for comparison. As shown in Tab. 1 (right) and Fig. 6, our method continues to outperform the alternatives. This comparison underscores that 2D deblurring methods fail to preserve the consistency of scene information. When the dynamic NeRFs are trained on these pre-processed images, it leads to undesirable artifacts in scene geometry. In contrast, our method models the blurring process in 3D space, ensuring spatial-temporal consistency of scene geometry.

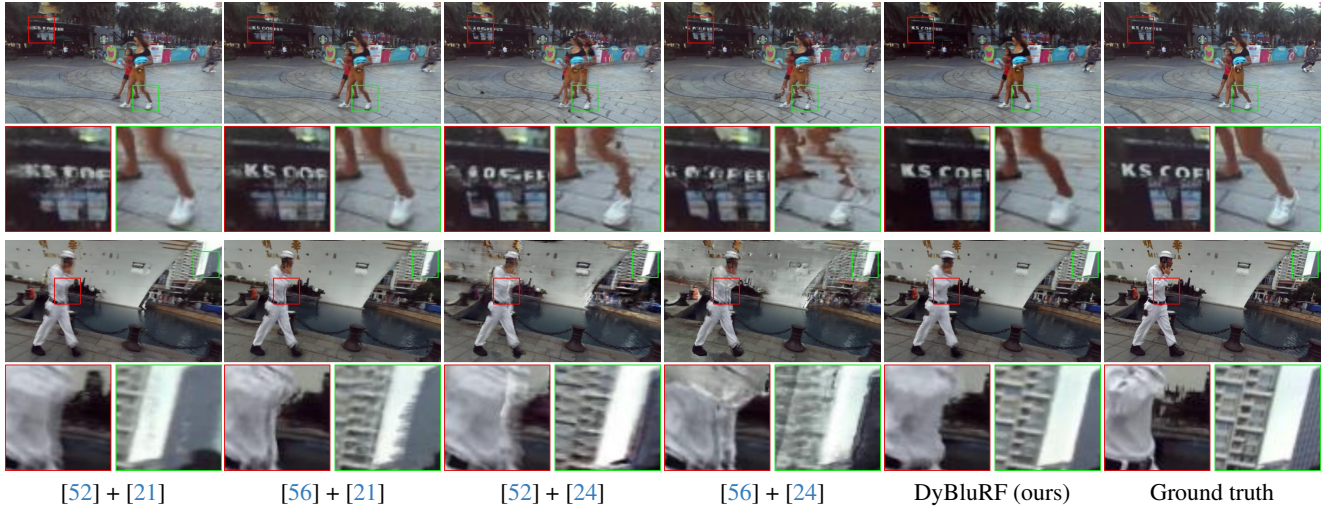


Figure 6. **Qualitative comparisons against dynamic NeRF baselines incorporated with 2D deblur method.** Even if we use preprocessed input blurry images by 2D deblur approaches to train existing dynamic NeRF methods, our method also generates more reliable novel views, with less blur in both static and dynamic regions.

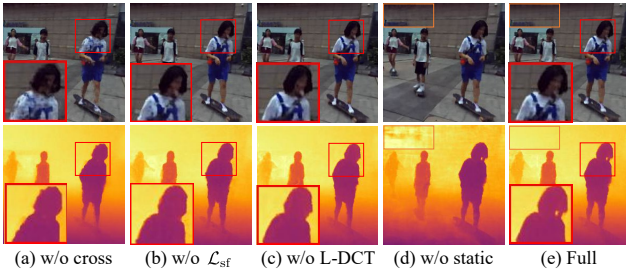


Figure 7. **Qualitative ablations.** (a) No cross-time rendering or (b) no \mathcal{L}_{sf} results in degradation of predicted RGB and depth maps, especially containing artifacts in dynamic regions. (c) Not learning DCT basis may impact the representation of scene details. (d) The absence of static model leads to poor results in static regions. (e) Our proposed full method yields high-quality results in both dynamic and static regions.

4.4. Ablation Study

We conduct an ablation study to evaluate the contribution of each component in our model. Specifically, we evaluate the impact of (1) removing the cross-time rendering (w/o cross), (2) using average in data-driven priors instead of EVC (w/o EVC), (3) removing the static model (w/o static), (4) removing \mathcal{L}_{sf} (w/o \mathcal{L}_{sf}), and (5) directly computing DCT basis using cosine values without being set as learnable parameters (w/o L-DCT). We report quantitative results in Tab. 2. One can see that without cross-time rendering and scene flow constraints, view synthesis quality degrades significantly. Combining static model can also improve performance. Furthermore, removing EVC or not learning DCT basis can reduce performance, demonstrating the value of these training tricks for improving novel view quality. We also conduct visual comparisons in Fig. 4(b) and Fig. 7, demonstrating that each component contributes

Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o cross	24.54	0.875	0.087
w/o EVC	25.42	0.890	0.079
w/o static	24.67	0.869	0.090
w/o \mathcal{L}_{sf}	24.97	0.879	0.085
w/o L-DCT	<u>25.49</u>	<u>0.891</u>	0.085
Full	25.66	0.895	0.079

Table 2. **Ablation study.** The best performance is **boldfaced**, and the second is underlined.

to the overall performance of the model. Specifically, without EVC in data-driven constraints, the incorrect depth and flow priors can misleadingly optimize the model to predict low-quality rendering results. Removing cross-time rendering or \mathcal{L}_{sf} cannot ensure the temporal consistency of dynamic scenes, resulting in obvious artifacts in dynamic regions. Static model can improve the scene representation for static elements, benefiting high-quality rendering results in static scenes. Finally, learning DCT basis can improve the representation of scene details.

5. Conclusion

We propose DyBluRF, a dynamic radiance field to synthesize sharp novel views from a monocular video with motion blur. By modeling camera trajectory and object motion trajectories within the scene, our method overcomes the impact of motion blur inputs on novel view synthesis. We also ensure the global temporal consistency of the predicted DCT trajectory through cross-time rendering. We conducted extensive experiments using scenes collected from the Stereo Blur Dataset, which demonstrate that our method outperforms existing dynamic NeRF methods in handling monocular inputs with motion blur.

References

- [1] Benjamin Attal, Jia-Bin Huang, Christian Richardt, Michael Zollhoefer, Johannes Kopf, Matthew O’Toole, and Changil Kim. Hyperreel: High-fidelity 6-dof video with ray-conditioned sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16610–16620, 2023. 3
- [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5855–5864, 2021. 2
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5470–5479, 2022. 2
- [4] Wenjing Bian, Zirui Wang, Kejie Li, Jia-Wang Bian, and Victor Adrian Prisacariu. Nope-nerf: Optimising neural radiance field with no pose prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4160–4169, 2023. 3
- [5] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 130–141, 2023. 3
- [6] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 333–350, 2022. 3
- [7] Zhiqin Chen, Thomas Funkhouser, Peter Hedman, and Andrea Tagliasacchi. Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16569–16578, 2023. 3
- [8] Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. In *ACM SIGGRAPH Asia*, pages 1–9, 2022. 3
- [9] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5501–5510, 2022. 3
- [10] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5712–5721, 2021. 1, 3, 5, 7
- [11] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18398–18408, 2022. 3
- [12] Tae Hyun Kim and Kyoung Mu Lee. Generalized video deblurring for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5426–5434, 2015. 2
- [13] James T Kajiya and Brian P Von Herzen. Ray tracing volume densities. *ACM SIGGRAPH computer graphics*, 18(3):165–174, 1984. 3
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [15] Dilip Krishnan, Terence Tay, and Rob Fergus. Blind deconvolution using a normalized sparsity measure. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 233–240, 2011. 2
- [16] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiří Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8183–8192, 2018. 2
- [17] Orest Kupyn, Tetiana Martyniuk, Junru Wu, and Zhangyang Wang. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8878–8887, 2019. 2
- [18] Dogyoon Lee, Minhyeok Lee, Chajin Shin, and Sangyoun Lee. Dp-nerf: Deblurred neural radiance field with physical scene priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12386–12396, 2023. 1, 3
- [19] Dongwoo Lee, Jeongtaek Oh, Jaesung Rim, Sunghyun Cho, and Kyoung Mu Lee. Exblurf: Efficient radiance fields for extreme motion blurred images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 17639–17648, 2023. 1, 3
- [20] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5521–5531, 2022. 3
- [21] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6498–6508, 2021. 1, 3, 4, 5, 6, 7, 8
- [22] Zhengqi Li, Qianqian Wang, Forrester Cole, Richard Tucker, and Noah Snavely. Dynibar: Neural dynamic image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4273–4284, 2023. 3, 5
- [23] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5741–5751, 2021. 3
- [24] Yu-Lun Liu, Chen Gao, Andreas Meuleman, Hung-Yu Tseng, Ayush Saraf, Changil Kim, Yung-Yu Chuang, Johannes Kopf, and Jia-Bin Huang. Robust dynamic radiance fields. In *Proceedings of the IEEE/CVF Conference on Com-*

- puter Vision and Pattern Recognition (CVPR), pages 13–23, 2023. 1, 3, 4, 5, 7, 8
- [25] Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12861–12870, 2022. 1, 3
- [26] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7210–7219, 2021. 3
- [27] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 405–421, 2020. 1, 2
- [28] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022. 3
- [29] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3883–3891, 2017. 2
- [30] Jinshan Pan, Haoran Bai, and Jinhui Tang. Cascaded deep video deblurring using temporal sharpness prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3043–3051, 2020. 2
- [31] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5865–5874, 2021. 3
- [32] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: a higher-dimensional representation for topologically varying neural radiance fields. *ACM Transactions on Graphics (TOG)*, 40(6):1–12, 2021. 1, 3, 7
- [33] Naama Pearl, Tali Treibitz, and Simon Korman. Nan: Noise-aware nerfs for burst-denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12672–12681, 2022. 3
- [34] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10318–10327, 2021. 1, 3
- [35] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)*, 44(3):1623–1637, 2020. 6
- [36] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4104–4113, 2016. 7
- [37] Qi Shan, Jiaya Jia, and Aseem Agarwala. High-quality motion deblurring from a single image. *ACM Transactions on Graphics (TOG)*, 27(3):1–10, 2008. 2
- [38] Ziyi Shen, Wenguan Wang, Xiankai Lu, Jianbing Shen, Haibin Ling, Tingfa Xu, and Ling Shao. Human-aware motion deblurring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5572–5581, 2019. 2
- [39] Shuo Chen, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1279–1288, 2017. 2
- [40] Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, and Jiaya Jia. Scale-recurrent network for deep image deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8174–8182, 2018. 2
- [41] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 402–419, 2020. 6
- [42] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12959–12970, 2021. 3
- [43] Jack Valmadre and Simon Lucey. General trajectory prior for non-rigid reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1394–1401, 2012. 2, 4
- [44] Chaoyang Wang, Ben Eckart, Simon Lucey, and Orazio Gallo. Neural trajectory fields for dynamic novel view synthesis. *arXiv preprint arXiv:2105.05994*, 2021. 3, 4
- [45] Chaoyang Wang, Lachlan Ewen MacDonald, Laszlo A Jeni, and Simon Lucey. Flow supervision for deformable nerf. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21128–21137, 2023. 3
- [46] Peng Wang, Lingzhe Zhao, Ruijie Ma, and Peidong Liu. Bad-nerf: Bundle adjusted deblur neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4170–4179, 2023. 1, 2, 3, 4
- [47] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021. 3
- [48] O. Whyte, J. Sivic, A. Zisserman, and J. Ponce. Non-uniform deblurring for shaken images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 491–498, 2010. 2
- [49] Wenqi Xian, Jia-Bin Huang, Johannes Kopf, and Changil Kim. Space-time neural irradiance fields for free-viewpoint

- video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9421–9431, 2021. [3](#)
- [50] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8254–8263, 2023. [3](#)
- [51] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4578–4587, 2021. [3](#)
- [52] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5728–5739, 2022. [2](#), [7](#), [8](#)
- [53] Huicong Zhang, Haozhe Xie, and Hongxun Yao. Spatio-temporal deformable attention network for video deblurring. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 581–596, 2022. [2](#)
- [54] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018. [7](#)
- [55] Zhihang Zhong, Ye Gao, Yinqiang Zheng, and Bo Zheng. Efficient spatio-temporal recurrent neural network for video deblurring. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 191–207, 2020. [2](#)
- [56] Zhihang Zhong, Yinqiang Zheng, and Imari Sato. Towards rolling shutter correction and deblurring in dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9219–9228, 2021. [7](#), [8](#)
- [57] Shangchen Zhou, Jiawei Zhang, Wangmeng Zuo, Haozhe Xie, Jinshan Pan, and Jimmy S Ren. Davanet: Stereo deblurring with view aggregation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10996–11005, 2019. [2](#), [6](#), [7](#)