

EventNeRF: Neural Radiance Fields from a Single Colour Event Camera

Viktor Rudnev^{1,2} Mohamed Elgharib¹ Christian Theobalt¹ Vladislav Golyanik¹

¹Max Planck Institute for Informatics, SIC ²Saarland University, SIC

Abstract

Asynchronously operating event cameras find many applications due to their high dynamic range, vanishingly low motion blur, low latency and low data bandwidth. The field saw remarkable progress during the last few years, and existing event-based 3D reconstruction approaches recover sparse point clouds of the scene. However, such sparsity is a limiting factor in many cases, especially in computer vision and graphics, that has not been addressed satisfactorily so far. Accordingly, this paper proposes the first approach for 3D-consistent, dense and photorealistic novel view synthesis using just a single colour event stream as input. At its core is a neural radiance field trained entirely in a self-supervised manner from events while preserving the original resolution of the colour event channels. Next, our ray sampling strategy is tailored to events and allows for data-efficient training. At test, our method produces results in the RGB space at unprecedented quality. We evaluate our method qualitatively and numerically on several challenging synthetic and real scenes and show that it produces significantly denser and more visually appealing renderings than the existing methods. We also demonstrate robustness in challenging scenarios with fast motion and under low lighting conditions. We release the newly recorded dataset and our source code to facilitate the research field, see <https://4dqv.mpi-inf.mpg.de/EventNeRF>.

1. Introduction

Event cameras record an asynchronous stream of per-pixel brightness changes. This is in contrast to conventional RGB cameras that record the absolute intensity values at a pre-defined framerate. Event cameras have several advantages over conventional cameras such as virtually no motion blur, higher dynamic range, ultra low power consumption and lower latency. Because of these advantages, event cameras have been applied to many computer vision problems, including optical-flow estimation [11, 12], video interpolation [26, 38, 57, 62], deblurring [26, 68], 3D pose estimation [36, 46, 64, 75], geometry reconstruction [1, 6, 58, 61]

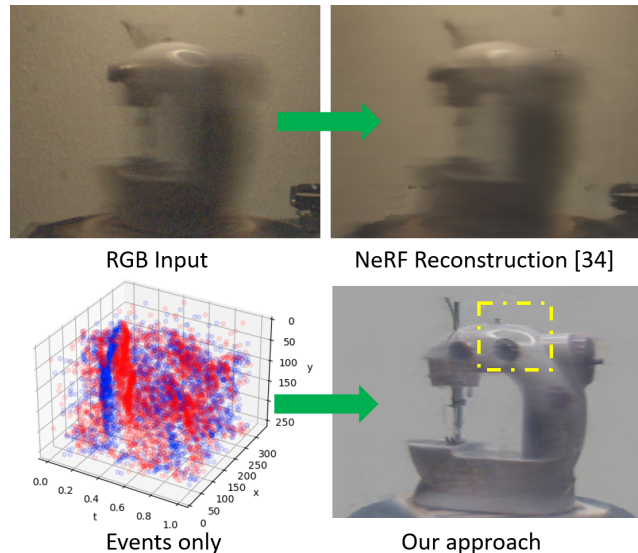


Figure 1. The original NeRF approach [35] (first row) produces strong artefacts when processing RGB images undergoing fast motion. Our approach (second row) operates purely on an event stream and thus produces significantly better and sharper results, capturing important details such as specularities (see yellow box).

and many more [4, 9, 51, 60, 62, 63, 71, 74].

Generating dense photorealistic rendering of a scene in a 3D-consistent manner is a long-standing problem in computer vision and computer graphics [7, 13, 23, 35, 49]. There is currently, however, no event-based method for solving this problem. The closest work in literature usually addresses the problem either from a SLAM [21, 41, 59, 72] or an odometry perspective [14, 22, 42, 73, 76]. Here, the aim is to estimate the position of a moving camera, and to reconstruct the 3D scene to some extent. These methods usually rely on explicit feature matching in the event space and reconstruct sparse 3D models of environments from a single event stream. As a result, novel views generated by existing methods are usually sparse, mostly containing edges and details causing events. The same observation applies to methods using stereo event cameras, even though they aim to densify the 3D reconstructions [72, 73].

The research question we are addressing in this work

is, whether *an event stream from an event camera moving around the scene is sufficient to reconstruct a dense volumetric 3D representation of a static scene*. Here, we adopt the state-of-the-art 3D representation of Neural Radiance Fields (NeRF). To this end, we present the first approach for inferring a NeRF volume from only a monocular colour event stream that enables 3D-consistent, dense and photorealistic novel view synthesis in the RGB space at test time. Our method, which we name *EventNeRF*, is designed for purely event-based supervision, during which we preserve the resolution of the individual RGB event channels. We evaluate our technique for the task of novel view synthesis on a new dataset of synthetic and real event sequences. This allows subjective and numerical evaluations against ground truth. We evaluate our NeRF estimation in scenarios that would not be conceivable with a traditional RGB camera (e.g., high speed movements, motion blur or insufficient lighting), and show our method produce significantly better results (see Fig. 1). We also show an application of extracting the depth map of an arbitrary viewpoint. To summarise, the primary **technical contributions** are as follows:

- 1) EventNeRF, the first approach for inferring NeRF from a monocular colour event stream that enable novel view synthesis in the RGB space at test time. Our method is designed for event-based supervision, during which we preserve the resolution of the individual RGB event channels (while avoiding demosaicing; see Sec. 3.4).
- 2) A ray sampling strategy tailored to events that allows for data-efficient training. As a part of it, our negative sampling avoids artefacts in the background (Sec. 3.7).
- 3) An experimental evaluation protocol for the task of novel view synthesis from event streams in the RGB space. We will release our code and dataset to establish a benchmark for future work (Sec. 4).

2. Related Work

Novel View Synthesis from Event Streams. The properties of event cameras (*i.e.*, no motion blur, high dynamic range, low latency and ultra-low power consumption) motivated their usage in computer vision [10, 12, 57, 60, 68]. While many of the developed techniques address problems from a 2D-vision perspective [9], event cameras have also been increasingly used for 3D-vision problems. This includes methods designed for dedicated objects such as the human body [64, 75] and the human hands [36, 46], as well as methods for general scenes [1, 6, 21, 41, 58, 59, 61]. Event-Cap [64] is a method for the 3D reconstruction of a human body in motion. Here, the tracked mesh can be used to render novel views of the performer. However, the method requires grayscale frames and a 3D body scan as inputs.

Thus, novel views of the actor do not contain new details observed in the input event streams (instead, the initial 3D scan is deforming). The analysis-by-synthesis method of Nevhi *et al.* [36] operates on events only and supports arbitrary non-rigid motions of the human hands. Rudnev *et al.* [46] regresses sparse 3D hand keypoints and drives a MANO [44] hand model from that. Even though the final results of both Nevhi *et al.* [36] or Rudnev *et al.* [36] can be in form of a mesh, none of them is designed for the task of photorealistic novel view synthesis.

Methods for the 3D reconstruction of general scenes have recently seen remarkable progress. They can be classified into ones that aim to estimate the underlying geometry only [1, 6, 58, 61] and others that also estimate the camera trajectory [21, 41, 59, 73, 76]. For the former, several methods have been proposed including ones that use a single event camera [1, 61], two event cameras [58] and other input data modalities such as LiDAR [6]. In the latter, methods follow either a SLAM [21, 41, 59, 72] or an odometry based formulation [14, 22, 42, 73, 76] and target large-scale scenes. In addition to reconstructing the 3D scene to some extent, these methods also estimate the position of a moving camera scanning the scene. To this end, some methods use as input a single event camera [21, 41, 42], a pair of event cameras [14, 72, 73] and some use a mix of input data modalities such as depth [76], intensity images [22] and IMUs [59]. While SLAM aims for denser 3D reconstruction of the underlying scene than odometry methods, both approach classes produce only sparse reconstructions (usually edges and corners). This is far from allowing 3D-consistent dense and photorealistic rendering of the examined scene. *In stark contrast to these works, we learn—for the first time—an implicit 3D scene representation from event streams that enables dense photorealistic novel view synthesis in the RGB space.*

3D Scene Representation Learning is a long-standing problem in computer vision and computer graphics [7, 13, 28, 34, 35, 49]. Some of the existing methods produce meshes [15, 19, 56] and multi-plane images (MPIs) [25, 34, 70]. Such representations can be learnt from 2D images [25, 55, 70], but they suffer from several limitations, including the inability to capture fine geometrical details (meshes) and being mostly bounded to 2.5D novel view synthesis (MPIs). In contrast, the recently introduced implicit coordinate-based representations learn to map the 3D spatial coordinates of a static scene to a continuous function of the local scene properties [28, 32, 35, 40, 49]. Here, Neural Radiance Fields, or NeRF [35], models a scene through MLP as a continuous function of the scene radiance and volume density learnt from a set of 2D RGB images. At test time, NeRF accepts arbitrary 3D camera pose and viewing direction to render novel views. Due to its simplicity and high accuracy of the generated novel views, multiple

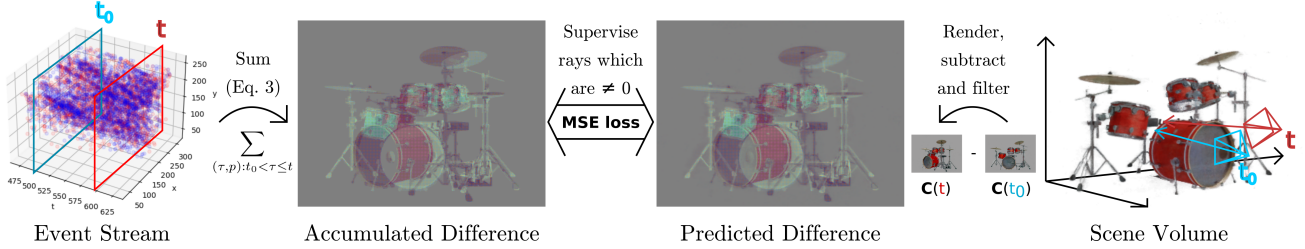


Figure 2. **Overview of the proposed method which learns a NeRF volume of a static object using just a moving event camera.** Our method connects the observed events $\{\mathbf{E}_i\}_{i=1}^N$ with the differences between rendered views (multiplied by Bayer colour filter) taken at two different time instants t_0 and t , via an event-based integral (see Sec. 3.2). The projections are estimated via a fully differentiable volumetric renderer (see Sec. 3.1), thus allowing learning the NeRF volume in a self-supervised manner.

follow-up works adopted NeRF for many applications such as reenactment [8, 27, 69], scene appearance and light editing [2, 30, 50], 3D shape modelling [3, 5, 16, 47] and others [37, 54, 65]. All these methods, however, assume input data captured by a traditional RGB camera.

3. Method

Our aim is to learn a neural 3D scene representation $\mathbf{C}(\mathbf{o}, \mathbf{d})$ for a static scene from just a colour event stream $\{\mathbf{E}_i\}_{i=1}^N$ (no RGB images are involved). The model $\mathbf{C}(\mathbf{o}, \mathbf{d})$ is a function of camera position $\mathbf{o} \in \mathbb{R}^3$ and its viewing direction $\mathbf{d} \in \mathbb{R}^3, \|\mathbf{d}\| = 1$, thus allowing novel view synthesis during test in the RGB space. To learn the model, we assume the camera extrinsics $\mathbf{P}_j = [\mathbf{R}_j | \mathbf{t}_j] \in \mathbb{R}^{3 \times 4}$ and intrinsics $\mathbf{K}_j \in \mathbb{R}^{3 \times 3}$ are known for each event window. We also assume a constant colour background with a known value $\alpha \in \mathbb{R}^3$ to rectify the reconstructed image brightness and colour balance. Its known value is not a strict requirement as the equivalent brightness and colour balance corrections can be done after the rendering. Our model is learned in a self-supervised manner, by comparing the approximate difference between views computed by summing the observed events polarities $\mathbf{E}(t_0, t)$ against the difference between the predicted views $\mathbf{L}(t) - \mathbf{L}(t_0)$ (Secs. 3.3, 3.4). To this end, we use NeRF [35] as our 3D scene representation, where each point in 3D-space \mathbf{x} is represented via a volume density $\sigma(\mathbf{x})$ and a radiance $\mathbf{c}(\mathbf{x})$ (Sec. 3.1). We then use volumetric rendering [35] to produce 2D projections of the learned $\mathbf{C}(\mathbf{o}, \mathbf{d})$, and establish a connection with the observed events via an event-based integral (Sec. 3.2). Since the volumetric rendering is fully differentiable, our model $\mathbf{C}(\mathbf{o}, \mathbf{d})$ can be learned in a completely self-supervised manner using events only. Our model is designed for rendering of novel view in the RGB space at test time; see Fig. 2. Lastly, we apply regularisation techniques and introduce an event-based ray sampling strategy that allows efficient learning of the model (Secs. 3.6-3.7).

3.1. NeRF and Volumetric Rendering

NeRF [35] learns a volumetric model of a static scene from multiple 2D RGB inputs captured at different camera viewpoints. It uses an MLP to map each point in 3D space into a continuous field of volume density and radiance. To render a 2D image, a ray is shot from a camera location along a specific direction, and observations are projected into 2D via volumetric rendering [20, 31].

Given a ray $\mathbf{r}(q) = \mathbf{o} + q\mathbf{d}$, where $\mathbf{o} \in \mathbb{R}^3$ is the camera origin and $\mathbf{d} \in \mathbb{R}^3, \|\mathbf{d}\| = 1$ is the ray direction, we want to render its colour $\mathbf{C}(\mathbf{r}) \in \mathbb{R}^3$. Random S points $\{\mathbf{x}_i\}_{i=1}^S, \mathbf{x}_i \in \mathbb{R}^3$ on the ray are sampled with corresponding depths $\{t_i\}_{i=1}^S, t_i \in [t_n, t_f]$. Here, t_n and t_f are the near and the far camera plane depths. The points $\{\mathbf{x}_i\}_{i=1}^S$ are encoded through positional encoding, akin to [48]: We feed the encoded representation $\gamma(x)$ into an MLP $f_1(\cdot)$ to obtain the point densities $\{\sigma_i\}_{i=1}^S$ and feature vectors $\{\mathbf{u}_i\}_{i=1}^S$ via $(\sigma_i, \mathbf{u}_i) = f_1(\gamma(\mathbf{x}_i))$. We then combine the result with the view direction \mathbf{d} into the final colours of the ray samples $\{\mathbf{c}_i\}_{i=1}^S$ using MLP $f_2(\cdot)$ as $\mathbf{c}_i = f_2(\sigma_i, \mathbf{u}_i, \gamma(\theta))$. Finally, $\{\mathbf{c}_i\}_{i=1}^S$ are integrated into the final colour $\mathbf{C}(\mathbf{r})$ using the volumetric rendering equation:

$$\mathbf{C}(\mathbf{r}) = \sum_{i=1}^S w_i \mathbf{c}_i, \text{ where } w_i = T_i(1 - \exp(-\sigma_i \kappa_i)), \quad (1)$$

$$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \kappa_j\right) \text{ and } \kappa_i = q_{i+1} - q_i.$$

We use hierarchical sampling as in the original NeRF but omit the corresponding notation.

3.2. Event-based Single Integral

We denote the absolute instantaneous intensity image at time t as $\mathbf{B}(t) \in \mathbb{R}^{W \times H}$ and its logarithmic image as $\mathbf{L}(t) = \frac{\log \mathbf{B}(t)}{g}$, where g a gamma correction value fixed to 2.2 in all experiments. Gamma correction is required to obtain linear colour from $\mathbf{B}(t)$ as it is intended for viewing

on the computer screen. It is encoded with sRGB gamma curve, and Gamma 2.2 is its recommended smooth approximation [53]. Each event is denoted as a tuple (t, x, y, p) , where t is the event timestamp, (x, y) are the 2D coordinates and $p \in \{-1, +1\}$ is the polarity. It is assumed to cause a $p\Delta$ change in the logarithmic absolute value $L_{x,y}(t)$ as compared to the value $L_{x,y}(t')$ at the timestamp t' when the previous event happened at (x, y) . Hence, $L_{x,y}(t) - L_{x,y}(t') = p\Delta$, where Δ is a fixed and polarity-wise symmetric event threshold. We model the change in $L_{x,y}(t)$ based on the events as follows:

$$L_{x,y}(t) - L_{x,y}(t_0) = \int_{t_0}^t \Delta e_{xy}(\tau) d\tau. \quad (2)$$

where t_0 and t are two different time instants (see Fig. 2). Here, $e_{xy}(\tau) = p\delta_\tau^t$ for all event tuples (t, x, y, p) , where δ is the impulse function with unit integral. By substituting $e_{xy}(\tau)$ into the integral, the following sum is obtained:

$$L_{x,y}(t) - L_{x,y}(t_0) = \sum_{(\tau,p):t_0 < \tau < t} p\Delta \stackrel{\text{def}}{=} E_{x,y}(t_0, t). \quad (3)$$

The right side of Eq. (3) is a constant computed from the event stream by accumulating per-pixel event polarities. The left part is the difference between the logarithmic absolute intensity values of the views rendered from $\mathbf{C}(\mathbf{o}, \mathbf{d})$ at t_0 and t . Thus, Eq. (3) establishes a link between the observed events $\{\mathbf{E}_i\}_{i=1}^N$ and the intensity images. We propose keeping the right side constant and model \mathbf{L} with the NeRF representation described in Sec. 3.1. We note that the described model (3) is similar to Event-based Double Integral (EDI) model proposed in Pan *et al.* [38]. However, with NeRF—as opposed to RGB cameras—we can render instantaneous moments, and, hence, there is no motion blur caused by the shutter of such cameras. Therefore, we reduce Double EDI to Single EDI in the model of Pan *et al.* [38].

3.3. Self-Supervision with Volumetric Rendering

We substitute the left hand side of Eq. (3) by the NeRF-based volumetric rendering \mathbf{C} from Eq. (1) described in Sec. 3.1. The differentiable nature of these renderings allows learning \mathbf{C} of the observed static scene in a self-supervised manner from a single event stream. We define our reconstruction loss function as the mean square error $\text{MSE}(\cdot)$ of the left side of (3) versus its right side:

$$\mathcal{L}_{\text{recon},xy}(t_0, t) = \text{MSE}(\hat{\mathbf{L}}_{xy}(t) - \hat{\mathbf{L}}_{xy}(t_0), E_{xy}(t_0, t)), \quad (4)$$

where $\hat{\mathbf{L}}_{xy}(t) = \frac{1}{g} \log \hat{\mathbf{B}}_{xy}(t)$, with g is a constant gamma correction as discussed earlier. As per Eq. (1), $\hat{\mathbf{B}}_{xy}(t) = \mathbf{C}(\mathbf{r}_{xy}(t))$, where $\mathbf{r}_{xy}(t)$ is the ray corresponding to the pixel at location (x, y) . Note that the derivations concern so far the grayscale model.

The right-hand argument of Eq. (4) is a scalar in case of a grayscale model. In this case, the MSE notation should be interpreted in the following way: The same grayscale value of the right argument is supervising *all three colour channels* of the left argument. This results in learning a grayscale $\hat{\mathbf{L}}(t)$. The supervision for the case of learning from colour events will be described next.

3.4. Using Colour Events

The difference between colour and grayscale event cameras is the presence of a colour (“Bayer”) filter in front of the sensor [43]. Hence, each pixel perceives one colour channel at the time. It is known and constant for each pixel and usually it is arranged in 2×2 blocks.

In traditional camera imaging, all missing colour channel values are recovered during *debayering* [24]. For each pixel, its neighbouring values are used to fill in the missing colour channels. In event cameras, events fire per pixel asynchronously and, hence, it is not possible to use traditional debayering. Another possible solution is to down-sample all 2×2 pixel blocks into a single pixel, at the cost of losing $3/4$ of spatial resolution. We next describe an approach to learn fully coloured model with neither maintaining full frame reconstructions nor losing spatial resolution.

We denote the Bayer filter as $\mathbf{F} \in \mathbb{R}^{W \times H \times 3}$, where $W \times H$ is the spatial resolution. Here, for each pixel (x, y) , $\mathbf{F}_{x,y,z} = 1$, only if z is the colour channel of the examined pixel, and 0 elsewhere. In the DAVIS 346C event camera which we use in our experiments, \mathbf{F} consists of the following tiled 2×2 pattern: $[[[1, 0, 0], [0, 1, 0]], [[0, 1, 0], [0, 0, 1]]]$ (RGGB colour filter).

To model the colour events, we point-wise pre-multiply the colour filter mask \mathbf{F} with both the rendered (left) and the event (right) arguments of the MSE in Eq. (4):

$$\mathcal{L}_{\text{recon}}(t_0, t) = \text{MSE}(\mathbf{F} \odot \hat{\mathbf{L}}(t) - \mathbf{F} \odot \hat{\mathbf{L}}(t_0), \mathbf{F} \odot E(t_0, t)), \quad (5)$$

where “ \odot ” denotes pixel-wise multiplication.

In pixels, where only red events happen, the multiplication by \mathbf{F} removes all other channels than red. This does not mean that a particular point in 3D space will only ever be supervised with by the red channel and have no green or blue channel information. Instead, when the same 3D point is observed from different views, it will be seen by green and blue pixels as well and, hence, have all the needed information to reconstruct the full RGB colour at the full resolution. Hence we can recover all three RGB channels without demosaicing or loss of spatial resolution of the event windows.

To learn the correct white balance of the scene, we make use of the assumption that the background colour $\alpha \in \mathbb{R}^3$ is known in advance. The object’s silhouette generates events capturing the difference between the background and foreground colours. This propagates the background colour information into the foreground object, thus recovering its

colour as well and allowing learning both the correct brightness offset and the colour balance between the channels.

3.5. Final Loss Function

Finally, we obtain the following loss function:

$$\mathcal{L} = \frac{1}{N_{\text{windows}}} \sum_{i=1}^{N_{\text{windows}}} \mathcal{L}_{\text{recon}}(t_0, t), \quad (6)$$

$$t = \frac{i}{N_{\text{windows}}}, t_0 \sim U[t - L_{\text{max}}, t),$$

where $N_{\text{windows}} = 1000$ in our experiments, and the choice of L_{max} is described in Sec. 3.7.

3.6. Ray Direction Jitter

An often-used camera trajectory is a circle around the object, so that the camera keeps roughly the same altitude and distance to the object. This makes it easy for the model to overfit to the input views. Following NeRF-OSR [45], we use ray direction jitter for better generalisation to unseen views, *i.e.*, instead of shooting the rays exactly through the pixel centre, we randomly offset the ray direction within the pixel. This policy is fast and easy to implement.

3.7. Event-based Ray Sampling Strategy

Our method can be used with individual events or event windows; we use event windows for efficiency and for simplicity. Thus, we accumulate events from the start to the end of the window. Now we describe how we select the starts and the ends of the windows. First, we split the whole event stream into intervals $t_i = i/N_{\text{windows}}$, where $i \in \{1, \dots, N_{\text{windows}}\}$; t_i are the window ends. We empirically noticed that using constant short windows results in poor propagation of high-level lighting; using constant long windows often results in poor local details. Hence, we sample the window length randomly as $t - t_0 \sim U(0, L_{\text{max}}]$. Equivalently, if the window end t is fixed, the window start is sampled as $t_0 \sim U[t - L_{\text{max}}, t)$. This allows the model to learn both global and local colour information. For real data experiments, we use $L_{\text{max}} = 0.1$, where $t = 1.0$ is the end of the stream. For synthetic data experiments, we use $L_{\text{max}} = 0.05$.

In every event window, N_e pixels are sampled from the positions where $E(t_0, t) \neq 0$, *i.e.*, at least one event happens and is not cancelled by others. However, if sampling is limited only by non-zero rays (*positive sampling*), information about pixels with no events will be ignored and would degrade results. Hence, we also apply *negative sampling*: We sample βN_e rays from every other pixel in the frame where $\beta = 0.1$. This makes our method more robust to noise events and leads to better reconstruction of the uniformly coloured surfaces as shown in the ablation studies (Sec. 4.4). It also helps in reconstructing thin structures.



Figure 3. Showing the fact that our real sequences were recorded under low lighting conditions. Here, we show the RGB stream of the event camera that was used for recording the real sequences. These sequences were recorded using just a 5W USB light source. We tone map the real sequences in the remaining figures for clarity.

In total, our sampling strategy uses only $O((1+\beta)N_e) = O(N_e)$ rays per epoch. This is much more efficient than the ray sampling strategy used in traditional NeRF for processing RGB images, where in each epoch, every pixel is rendered from every camera viewpoint, thus requiring $O(W \times H \times N_{\text{views}})$ rays, where $W \times H$ is the resolution and N_{views} the number of camera viewpoints. This observation proposes to discard the traditional notion of training view count [54] in our work and to instead use new data-efficiency metrics based on the number of events. It also suggests that redundant views do not help model learning as much as when they were diversely sampled.

3.8. Implementation Details

Our code is based on NeRF++ [66] without the background network as our test scenes can completely fit inside the unit sphere rendered using the foreground network. We replace ReLU with tanh activation function as it results in more consistent training results. We train the models for $5 \cdot 10^5$ iterations on two NVIDIA Quadro RTX 8000 GPUs with the optimiser and hyper-parameters as in NeRF++. This takes about six hours to complete. Rendering a full 346×260 image (*i.e.*, the resolution of DAVIS 346C) then takes about 1.7 seconds using the same hardware setup. Our real-time demo is based on torch-ngp [52] (see the supplement). We train a model for $5 \cdot 10^3$ iterations on a single NVIDIA RTX 2070, which takes around one minute to converge. After training, rendering in Full HD resolution is real-time at 30 FPS using the same hardware setup.

4. Experiments

Examined Sequences. We examine seven synthetic and ten real sequences. As synthetic ones, we use the 3D models from Mildenhall *et al.* [35]. For each scene, we render a one-second-long 360° rotation of camera around the object at 1000 fps as RGB images, resulting in 1000 views. From these images, we simulate the events stream using the model from [46]. The corresponding camera intrinsics and extrin-



Figure 4. Reconstructions by our EventNeRF on synthetic (“Materials”, “Ficus”) and real (“Goatling”, “Sewing”) sequences. Our approach reconstructs specularities (“Materials”), thin structure (“Ficus”) and even the eyeglass frame and the needle from the “Goatling” and “Sewing” sequences respectively (see yellow arrow).

sics are used directly in our method. For the real-data experiments, we record ten objects with the DAVIS 346C colour event camera on a uniform white background. We notice that it is hard to make a stable and calibrated setup in real life where the event camera rotates around an object. In the absence of the background events, rotating the object while keeping the camera still results in the same event stream as when the camera rotates around the static object. Hence, we keep the camera static and place the objects on a direct drive vinyl turntable, resulting in stable and consistent object rotation at 45 RPM. In this setting, it is essential to provide constant lighting regardless of the current object’s rotation angle, and we mount a single USB ring light above the object. A photo of the setup is included in the supplement (Fig. 1). Please note that all the real sequences are captured under low lighting condition (see Fig. 3), as the light source we use have a power of just 5W.

Metrics. To account for the varying event generation thresholds, colour balance, and exposure, we do the following procedure for all our and baseline results. In our model, the event threshold only affects the overall image contrast in logarithmic space. Exposure and colour balance are modelled as an offset in logarithmic space. Hence for every sequence of predicted images \mathbf{I}_k and the corresponding ground-truth images \mathbf{G}_k , we fit a single linear colour transform $f(\mathbf{I})$, which we apply to the predictions in the logarithmic space:

$$f(\mathbf{I}) = \exp(\mathbf{a} \odot \log \mathbf{I} + \mathbf{b}),$$

$$(\mathbf{a}, \mathbf{b}) = \arg \min_{\mathbf{a}, \mathbf{b} \in \mathbb{R}^3} \sum_{k=1}^{N_{\text{images}}} \|\mathbf{a} \odot \log \mathbf{I}_k + \mathbf{b} - \log \mathbf{G}_k\|^2. \quad (7)$$

Then, for the transformed images $f(\mathbf{I}_k)$, we report PSNR, SSIM, and LPIPS [67] based on AlexNet.

4.1. Synthetic Sequences

As described earlier, we examine synthetic sequences from Mildenhall *et al.* [35]. They cover different effects such as thin structures (drums, caterpillar, ficus, microphone), view-dependent effects (drums, caterpillar, materials), large uniformly coloured surfaces (chair). Fig. 4-(left) shows visual results for two sequences. EventNeRF learns view-dependent effects (Materials) and thin structures (Ficus). Fig. 5 shows our reconstruction on two more synthetic sequences, Here, we capture textured regions well (Lego). Corresponding numerical results are reported in Tab. 1.

4.2. Real Sequences

We examine ten scenes recorded on a turntable: Goatling, dragon, chicken, sewing machine, game controller, cube, bottle, fake plant, multimeter and microphone. They show the performance of our method on scenes with thin structures (goatling, sewing machine, fake plant, microphone), fine-print coloured text (game controller, multimeter, bottle), view-dependent effects (goatling, sewing machine, cube, bottle), dark details (goatling, game controller). As there is no ground-truth RGB data, we can only evaluate results visually, as shown in Figs. 4 and 6. Note how can we can reconstruct a one-pixel-wide sewing needle in the “Sewing” sequence, and the eyeglasses frame on the “Goatling” scene (see yellow arrows in Fig. 4). Our results are also halo-free despite being reconstructed from events only. In Fig. 8, we show the extracted depth. Here, we use a colour scheme that represents per-pixel distances in meters from a virtual camera to the object in a novel view.




Figure 5. EventNeRF better resembles the ground truth than the baseline E2VID+NeRF.

Scene	E2VID [43] + NeRF [35]			Our EventNeRF		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Drums	19.71	0.85	0.22	27.43	0.91	0.07
Lego	20.17	0.82	0.24	25.84	0.89	0.13
Chair	24.12	0.92	0.12	30.62	0.94	0.05
Ficus	24.97	0.92	0.10	31.94	0.94	0.05
Mic	23.08	0.94	0.09	31.78	0.96	0.03
Hotdog	24.38	0.93	0.12	30.26	0.94	0.04
Materials	22.01	0.92	0.13	24.10	0.94	0.07
Average	22.64	0.90	0.15	28.85	0.93	0.06

Table 1. Comparing our method against E2VID+NeRF. Our method consistently produces better results.

4.3. Comparisons against Related Methods

We first recover RGB frames from events using E2VID [43] and then learn original frame-based NeRF with them as the training data; see Tab. 1. EventNeRF clearly outperforms this approach (which we call E2VID+NeRF) in all metrics and on all sequences. Fig. 5 shows visual samples of this experiment on the Drums and Lego sequences. Fig. 6 show that E2VID+NeRF also generates noticeable artefacts on the real sequences. Note that replacing E2VID with ssl-E2VID [39] generates even worse results, as the performance of ssl-E2VID is usually bounded by E2VID (see the supplement). Note that such approaches do not account for the sparsity and asynchronicity of the event stream and need much more memory and disk space to store all the reconstructed views. Moreover, there is a limit to how short the E2VID window can be made, and hence to the number of reconstructed views. In contrast, EventNeRF respects the asynchronous nature of event streams and reconstructs the neural representation directly from them; it can use an arbitrary number of windows, which allows reconstructing the scene even from 3% of the data, as we show in Sec. 4.4. That makes it significantly more scalable than first reconstructing the frames and using traditional NeRF [35].




Figure 6. On real data our EventNeRF clearly outperforms all the related methods of E2VID+NeRF and Deblur-NeRF.

We also compare against Deblur-NeRF [29], *i.e.*, a RGB-based NeRF extraction method designed specifically to handle blurry RGB videos. We perform this comparison by applying Deblur-NeRF on the the RGB stream of the event camera. Results in Fig. 6, however, show that our approach significantly outperform Deblur-NeRF. This is expected, as Deblur-NeRF can only handle view-inconsistent blur. On the other hand, our approach produces almost blur-free results. In addition, it is significantly more memory and computationally efficient. For Deblur-NeRF, 100 training views were used, which took around 22 seconds to record due to the low-light condition (see Fig. 3). Our EventNeRF approach, however, need only one 1.33s revolution of the object. In addition, it converges well within 6 hours using two NVIDIA Quadro RTX 8000 GPUs. This compares favourably to Deblur-NeRF, which needs 16 hours with the same GPU resources. Furthermore, the training time for our method significantly drops with the torch-ngp [52] implementation to just 1 minute using a single NVIDIA GeForce RTX 2070 GPU. This, however, comes with some compromise in the rendering quality (see the supplemental video). We also evaluated Deblur-NeRF on two synthetic sequences where we simulated the blur to match real recorded sequences. Our technique performs significantly better in PSNR, SSIM and LPIPS: (27.43, 0.91, 0.07) for Drums and (25.84, 0.89, 0.13) for Lego; and Deblur-NeRF scores much worse at (21.61, 0.76, 0.36) and (21.06, 0.76, 0.35).




Figure 7. Importance of our various design choices (see Sec. 4.4). The full model produces the best results.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Fixed 50 ms win.	27.32	0.90	0.09
W/o neg. simpl.	26.48	0.87	0.16
Full EventNeRF	27.43	0.91	0.07

Table 2. Ablation studies computed on the *Drums*.

Our results on real data show that EventNeRF handles poorly lit environments (all our sequences were shot in dim lighting; see Fig. 3). Note that we cannot compare against RAW-NeRF [33] and HDR-NeRF [17] as they require fundamentally different setups. RAW-NeRF [33] requires the RAW data as input and HDR-NeRF [17] requires footage of the same object shot with different exposures. Hence, they cannot be applied to the RGB frames of an event camera.

4.4. Ablation Study and Data Efficiency

We ablate the use of negative sampling and analyse window size randomisation as an alternative to a fixed long window. The results are shown both in Tab. 2 and in Fig. 7 on “Drums”. The most significant degradation comes from disabling negative sampling, as this leads to the highly visible artefacts in the background and from the object. Furthermore, using fixed 50 ms window results in the loss of fine detail and the emergence of halo in the reconstruction.

By varying the event generation threshold Δ in the simulated event streams, we can vary the total number of events. With higher Δ , the simulated camera becomes less sensitive to brightness changes in a fashion similar to quantisation. At a high enough Δ , some details in the scene never trigger a single event. Hence, we can compare how our method performs with different event numbers. In the Fig. 9, we plot




Figure 8. Examples of depth maps (in meters) at arbitrary novel views obtained in our tests with EventNeRF with real data (left: “Sewing”, right: “Goatling”). Note the fine details such as the sewing needle and the eyeglasses frame highlighted by the arrows. Rays that do not hit the object have zero depth.

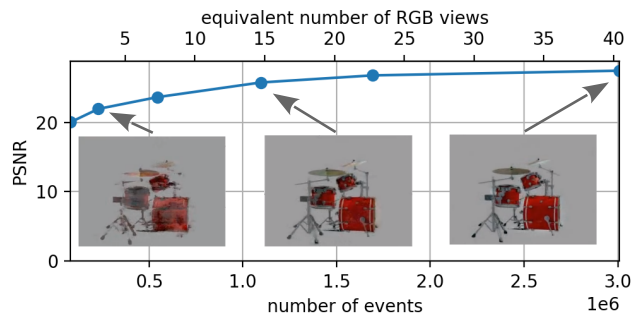


Figure 9. PSNR performance of our EventNeRF on “Drums” with varying event numbers (multiplied by 10^6 on the x -axis). The images underneath the function graph show exemplary novel views obtained with the 7%, 36% and 100% of the events, respectively.

PSNR of the reconstructed “Drums” scene as a function of the number of events and the number of equivalent RGB frames (that would occupy the same space). The generated novel views with 36% and 100% of the available events are barely distinguishable and with 7% of the data, the scene is still well recognisable. Note how we can still reconstruct the scene with the amount of data that is less than **one** equivalent RGB frame at 20 PSNR. For reference, NeRF’s performance degrades significantly with less than 100 views [18].

5. Conclusion

We introduced the first method to reconstruct a 3D model of a static scene from event streams that enables dense photorealistic RGB view synthesis. Thanks to the proposed combination of event supervision, volumetric rendering and several event-specific regularising techniques, EventNeRF outperforms the baselines in the rendered image quality, fast motion handling, low-illumination handling, and data storage requirements. Thus, this paper extends the spectrum of practical event-based techniques with a 3D representation learning approach. Future work can investigate joint estimation of the camera parameters and the NeRF volume.

Acknowledgements. This work was supported by the ERC Consolidator Grant *4DReply* (770784).

References

- [1] Alexis Baudron, Zihao W Wang, Oliver Cossairt, and Aggelos K Katsaggelos. E3d: event-based 3d shape reconstruction. *arXiv preprint arXiv:2012.05214*, 2020. 1, 2
- [2] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P.A. Lensch. Nerd: Neural reflectance decomposition from image collections. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [3] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Computer Vision and Pattern Recognition (CVPR)*, pages 5799–5809, 2021. 3
- [4] Zehao Chen, Qian Zheng, Peisong Niu, Huajin Tang, and Gang Pan. Indoor lighting estimation using an event camera. In *Computer Vision and Pattern Recognition (CVPR)*, pages 14755–14765, 2021. 1
- [5] Enric Corona, Tomas Hodan, Minh Vo, Francesc Moreno-Noguer, Chris Sweeney, Richard Newcombe, and Lingni Ma. Lisa: Learning implicit shape and appearance of hands. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [6] Mingyue Cui, Yuzhang Zhu, Yechang Liu, Yunchao Liu, Gang Chen, and Kai Huang. Dense depth-map estimation based on fusion of event camera and sparse lidar. *Transactions on Instrumentation and Measurement (TIM)*, 71:1–11, 2022. 1, 2
- [7] Paul E. Debevec, Camillo J. Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach. In *SIGGRAPH*, 1996. 1, 2
- [8] Guy Gafni, Justus Thies, Michael Zollhöfer, and Matthias Nießner. Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In *Computer Vision and Pattern Recognition (CVPR)*, pages 8649–8658, 2021. 3
- [9] Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Tabbara, Andrea Censi, Stefan Leutenegger, Andrew J. Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 44(01):154–180, 2022. 1, 2
- [10] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. EKLt: Asynchronous, photometric feature tracking using events and frames. *International Journal on Computer Vision (IJCV)*, 2019. 2
- [11] Mathias Gehrig, Mario Millhäusler, Daniel Gehrig, and Davide Scaramuzza. E-raft: Dense optical flow from event cameras. In *International Conference on 3D Vision (3DV)*, 2021. 1
- [12] Mathias Gehrig, Manasi Muglikar, and Davide Scaramuzza. Dense continuous-time optical flow from events and frames. *arXiv preprint arXiv:2203.13674*, 2022. 1, 2
- [13] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. In *SIGGRAPH*, page 43–54, 1996. 1, 2
- [14] Antea Hadviger, Igor Cvišić, Ivan Marković, Sacha Vražić, and Ivan Petrović. Feature-based event stereo visual odometry. In *European Conference on Mobile Robots (ECMR)*, pages 1–6. IEEE, 2021. 1, 2
- [15] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 37(6):257:1–257:15, 2018. 2
- [16] Yang Hong, Bo Peng, Haiyao Xiao, Ligang Liu, and Juyong Zhang. Headnerf: A real-time nerf-based parametric head model. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [17] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High dynamic range neural radiance fields. In *Computer Vision and Pattern Recognition (CVPR)*, pages 18398–18408, 2022. 8
- [18] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *International Conference on Computer Vision (ICCV)*, pages 5885–5894, 2021. 8
- [19] Navami Kairanda, Edith Tretschk, Mohamed Elgharib, Christian Theobalt, and Vladislav Golyanik. ϕ -SfT: Shape-from-template with a physics-based deformation model. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [20] James T. Kajiya and Brian P Von Herzen. Ray tracing volume densities. *ACM Transactions on Graphics*, 18(3):165–174, 1994. 3
- [21] Hanme Kim, Stefan Leutenegger, and Andrew J. Davison. Real-time 3d reconstruction and 6-dof tracking with an event camera. In *European Conference on Computer Vision (ECCV)*, 2016. 1, 2
- [22] Beat Kueng, Elias Mueggler, Guillermo Gallego, and Davide Scaramuzza. Low-latency visual odometry using event-based feature tracks. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 16–23, 2016. 1, 2
- [23] Marc Levoy and Pat Hanrahan. Light field rendering. In *23rd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '96*, page 31–42, New York, NY, USA, 1996. Association for Computing Machinery. 1
- [24] Xin Li, Bahadır Gunturk, and Lei Zhang. Image demosaicing: A systematic survey. In *Visual Communications and Image Processing (VCIP)*, volume 6822, pages 489–503. SPIE, 2008. 4
- [25] Zhengqi Li, Wenqi Xian, Abe Davis, and Noah Snavely. Crowdsampling the plenoptic function. In *European Conference on Computer Vision (ECCV)*, pages 178–196. Springer, 2020. 2
- [26] Songnan Lin, Jiawei Zhang, Jinshan Pan, Zhe Jiang, Dongqing Zou, Yongtian Wang, Jing Chen, and Jimmy Ren. Learning event-driven video deblurring and interpolation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *European Conference on Computer Vision (ECCV)*, pages 695–710, Cham, 2020. Springer International Publishing. 1
- [27] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose con-

- trol. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)*, 2021. 3
- [28] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Transaction on Graphics*, 38(4):65:1–65:14, July 2019. 2
- [29] 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 *Computer Vision and Pattern Recognition (CVPR)*, pages 12861–12870, 2022. 7
- [30] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [31] Nelson Max. Optical models for direct volume rendering. *Transactions on Visualization and Computer Graphics (TVCG)*, 1(2):99–108, 1995. 3
- [32] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [33] Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P Srinivasan, and Jonathan T Barron. Nerf in the dark: High dynamic range view synthesis from noisy raw images. In *Computer Vision and Pattern Recognition (CVPR)*, pages 16190–16199, 2022. 8
- [34] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics*, 38(4), 2019. 2
- [35] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 3, 5, 6, 7
- [36] Jalees Nehvi, Vladislav Golyanik, Franziska Mueller, Hans-Peter Seidel, Mohamed Elgharib, and Christian Theobalt. Differentiable event stream simulator for non-rigid 3d tracking. In *CVPR Workshop on Event-based Vision*, 2021. 1, 2
- [37] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [38] Liyuan Pan, Cedric Scheerlinck, Xin Yu, Richard Hartley, Miaomiao Liu, and Yuchao Dai. Bringing a blurry frame alive at high frame-rate with an event camera. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6820–6829, 2019. 1, 4
- [39] Federico Paredes-Vallés and Guido C. H. E. de Croon. Back to event basics: Self-supervised learning of image reconstruction for event cameras via photometric constancy. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3445–3454, 2021. 7
- [40] Jeong Joon Park, Peter R. Florence, Julian Straub, Richard A. Newcombe, and S. Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. *Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [41] Henri Rebecq, Guillermo Gallego, Elias Mueggler, and Davide Scaramuzza. Emvs: Event-based multi-view stereo–3d reconstruction with an event camera in real-time. *Int. J. Comput. Vision (IJCV)*, 126(12):1394–1414, 2018. 1, 2
- [42] Henri Rebecq, Timo Horstschaefer, Guillermo Gallego, and Davide Scaramuzza. Evo: A geometric approach to event-based 6-dof parallel tracking and mapping in real time. *Robotics and Automation Letters (RA-L)*, 2(2):593–600, 2017. 1, 2
- [43] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. High speed and high dynamic range video with an event camera. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2019. 4, 7
- [44] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6), Nov. 2017. 2
- [45] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. Nerf for outdoor scene relighting. In *European Conference on Computer Vision (ECCV), Part XVI*, pages 615–631. Springer, 2022. 5
- [46] Viktor Rudnev, Vladislav Golyanik, Jiayi Wang, Hans-Peter Seidel, Franziska Mueller, Mohamed Elgharib, and Christian Theobalt. Eventhands: Real-time neural 3d hand pose estimation from an event stream. In *International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 5
- [47] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [48] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [49] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 1, 2
- [50] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [51] Timo Stoffregen, Guillermo Gallego, Tom Drummond, Lindsay Kleeman, and Davide Scaramuzza. Event-based motion segmentation by motion compensation. In *International Conference on Computer Vision (ICCV)*, pages 7244–7253, 2019. 1

- [52] Jiayang Tang. Torch-ngp: a pytorch implementation of instant-ngp, 2022. <https://github.com/ashawkey/torch-ngp>. 5, 7
- [53] TC 100/TA 2. International Standard IEC 61966-2-1:1999: Amendment 1 - Multimedia systems and equipment – Colour measurement and management – Part 2-1: Colour management – Default RGB colour space – sRGB. Standard, International Electrotechnical Commission, 2003. 4
- [54] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edith Treitsch, Wang Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, Tomas Simon, Christian Theobalt, Matthias Nießner, Jonathan T. Barron, Gordon Wetzstein, Michael Zollhöfer, and Vladislav Golyanik. Advances in Neural Rendering. *Computer Graphics Forum (EG STAR)*, 2022. 3, 5
- [55] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics*, 38(4), 2019. 2
- [56] Justus Thies, Michael Zollhöfer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2Face: Real-time Face Capture and Reenactment of RGB Videos. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [57] Stepan Tulyakov, Daniel Gehrig, Stamatios Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza. TimeLens: Event-based video frame interpolation. *Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2
- [58] S. M. Nadim Uddin, Soikat Hasan Ahmed, and Yong Ju Jung. Unsupervised deep event stereo for depth estimation. *Transactions on Circuits and Systems for Video Technology (TCSVT)*, 32(11):7489–7504, 2022. 1, 2
- [59] Antoni Rosinol Vidal, Henri Rebecq, Timo Horstschaef, and Davide Scaramuzza. Ultimate slam? combining events, images, and imu for robust visual slam in hdr and high-speed scenarios. *Robotics and Automation Letters (RA-L)*, 3(2):994–1001, 2018. 1, 2
- [60] Lin Wang, Tae-Kyun Kim, and Kuk-Jin Yoon. Eventsr: From asynchronous events to image reconstruction, restoration, and super-resolution via end-to-end adversarial learning. In *Computer Vision and Pattern Recognition (CVPR)*, pages 8312–8322, 2020. 1, 2
- [61] Ziyun Wang, Kenneth Chaney, and Kostaexacts Daniilidis. Evac3d: From event-based apparent contours to 3d models via continuous visual hulls. In *European conference on computer vision*, 2022. 1, 2
- [62] Ziwei Wang, Yonhon Ng, Cedric Scheerlinck, and Robert Mahony. An asynchronous kalman filter for hybrid event cameras. In *International Conference on Computer Vision (ICCV)*, pages 448–457, October 2021. 1
- [63] Kun Xiao, Guohui Wang, Yi Chen, Jinghong Nan, and Yongfeng Xie. Event-based dense reconstruction pipeline. In *International Conference on Robotics and Automation Sciences (ICRAS)*, pages 172–177, 2022. 1
- [64] Lan Xu, Weipeng Xu, Vladislav Golyanik, Marc Habermann, Lu Fang, and Christian Theobalt. Eventcap: Monocular 3d capture of high-speed human motions using an event camera. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2
- [65] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [66] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 5
- [67] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [68] Xiang Zhang and Lei Yu. Unifying motion deblurring and frame interpolation with events. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2
- [69] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C. Bühler, Xu Chen, Michael J. Black, and Otmar Hilliges. I M Avatar: Implicit morphable head avatars from videos. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [70] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. In *SIGGRAPH*, 2018. 2
- [71] Yi Zhou, Guillermo Gallego, Xiuyuan Lu, Siqi Liu, and Shaojie Shen. Event-based motion segmentation with spatio-temporal graph cuts. *Transactions on Neural Networks and Learning Systems (TNNLS)*, 2021. 1
- [72] Yi Zhou, Guillermo Gallego, Henri Rebecq, Laurent Kneip, Hongdong Li, and Davide Scaramuzza. Semi-dense 3d reconstruction with a stereo event camera. In *European Conference on Computer Vision (ECCV)*, pages 235–251, 2018. 1, 2
- [73] Yi Zhou, Guillermo Gallego, and Shaojie Shen. Event-based stereo visual odometry. *Transactions on Robotics (T-RO)*, 37(5):1433–1450, 2021. 1, 2
- [74] Dekai Zhu, Zhongcong Xu, Jinhu Dong, Canbo Ye, Yinbai Hu, Hang Su, Zhengfa Liu, and Guang Chen. Neuromorphic visual odometry system for intelligent vehicle application with bio-inspired vision sensor. In *International Conference on Robotics and Biomimetics (ROBIO)*, pages 2225–2232, 2019. 1
- [75] Shihao Zou, Chuan Guo, Xinxin Zuo, Sen Wang, Hu Xiaoqin, Shoushun Chen, Minglun Gong, and Li Cheng. Event-thpe: Event-based 3d human pose and shape estimation. In *International Conference on Computer Vision (ICCV)*, 2021. 1, 2
- [76] Yi-Fan Zuo, Jiaqi Yang, Jiaben Chen, Xia Wang, Yifu Wang, and Laurent Kneip. Devo: Depth-event camera visual odometry in challenging conditions. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2179–2185, 2022. 1, 2