

Robust *e*-NeRF: NeRF from Sparse & Noisy Events under Non-Uniform Motion

Weng Fei Low Gim Hee Lee

The NUS Graduate School's Integrative Sciences and Engineering Programme (ISEP)
Institute of Data Science (IDS), National University of Singapore
Department of Computer Science, National University of Singapore

{wengfei.low, gimhee.lee}@comp.nus.edu.sg

<https://wengflow.github.io/robust-e-nerf>

Abstract

Event cameras offer many advantages over standard cameras due to their distinctive principle of operation: low power, low latency, high temporal resolution and high dynamic range. Nonetheless, the success of many downstream visual applications also hinges on an efficient and effective scene representation, where Neural Radiance Field (NeRF) is seen as the leading candidate. Such promise and potential of event cameras and NeRF inspired recent works to investigate on the reconstruction of NeRF from moving event cameras. However, these works are mainly limited in terms of the dependence on dense and low-noise event streams, as well as generalization to arbitrary contrast threshold values and camera speed profiles. In this work, we propose Robust *e*-NeRF, a novel method to directly and robustly reconstruct NeRFs from moving event cameras under various real-world conditions, especially from sparse and noisy events generated under non-uniform motion. It consists of two key components: a realistic event generation model that accounts for various intrinsic parameters (e.g. time-independent, asymmetric threshold and refractory period) and non-idealities (e.g. pixel-to-pixel threshold variation), as well as a complementary pair of normalized reconstruction losses that can effectively generalize to arbitrary speed profiles and intrinsic parameter values without such prior knowledge. Experiments on real and novel realistically simulated sequences verify our effectiveness. Our code, synthetic dataset and improved event simulator are public.

1. Introduction

Event cameras are bio-inspired sensors that represent a paradigm shift in visual acquisition and processing. This is attributed to its fundamentally distinctive principle of operation, where its pixels independently respond to log-intensity changes in an asynchronous manner, yielding a stream of *events*, rather than measuring absolute linear intensity syn-

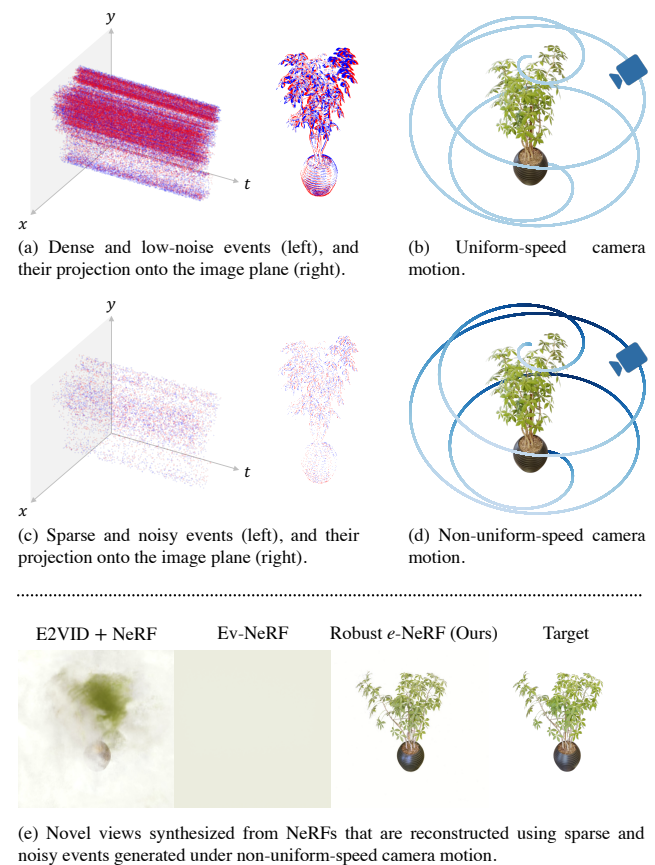


Figure 1. Existing works on NeRF reconstruction from moving event cameras heavily rely on (a) temporally dense and low-noise events generated under roughly (b) uniform-speed camera motion. In contrast, our method, Robust *e*-NeRF is able to directly and robustly reconstruct NeRFs from (c) sparse and noisy events generated under (d) non-uniform camera motion, as shown in (e).

chronously at a constant rate, as done in standard cameras. Such unique properties contribute to their multitude of advantages over standard cameras [9]: *low power, low la-*

gency, high temporal resolution and high dynamic range, thereby the recent success of event-based [15, 37, 10, 7, 38] or event-image hybrid [48, 47, 11] applications.

The success of many downstream visual applications in robotics, computer vision, graphics and virtual/augmented reality also hinges on an efficient and effective representation to encode various information of the scene being interacted with. Neural scene representations [24, 28, 49, 54, 20], especially neural fields [52], have recently emerged as promising candidates for future applications, owing to their continuous nature and memory efficiency. This trend is further driven by the exceptional capabilities and photorealism of *Neural Radiance Field* (NeRF) [24]-based works [49, 23, 45, 30, 31, 32].

Motivated by such promise and potential of event cameras and NeRF, we are interested in studying the following research question: *How to robustly reconstruct a NeRF from a moving event camera under general real-world conditions?* One simple way is to retrofit an events-to-video reconstruction method [38, 44, 29] to NeRF. However, such a naïve approach is inherently limited by the low photometric accuracy and consistency of the reconstructed video frames, since they are assumed to be true observations of the scene.

On the contrary, recent works, such as EventNeRF [41], Ev-NeRF [12] and E-NeRF [18], have proposed to reconstruct NeRFs directly using events via alternative reconstruction losses inspired or derived from an event generation model. Nonetheless, these works heavily rely on a temporally dense and low-noise event stream, which is generally inaccessible in practice due to the presence of *refractory period* (i.e. pixel dead-time after generating an event) and pixel-to-pixel variation in the *contrast threshold* (i.e. minimum log-intensity change for event generation). Such a limitation can be partly attributed to the accumulation of successive events at each pixel over time intervals, as performed in these works. Moreover, the effective reduction in contrast sensitivity and amplification of threshold variation resulting from the event accumulation also leads to a less detailed and robust reconstruction, respectively.

In addition, these methods do not directly and effectively generalize to arbitrary contrast threshold values and camera speed profiles, as their optimal hyper-parameter configuration greatly depends on the contrast threshold and speed of motion. EventNeRF and E-NeRF assumes symmetric positive and negative thresholds, which generally does not hold true in practice. While joint optimization of the contrast threshold is supported in Ev-NeRF, an additional regularization is necessary to prevent degeneracy. Furthermore, the assumption of time-varying thresholds made in Ev-NeRF and E-NeRF, which is not well supported by the literature, also leads to a reduction in reconstruction accuracy as shown in our experiments.

Contributions. We propose Robust *e*-NeRF, a novel method to directly and robustly reconstruct NeRFs from moving event cameras under various real-world conditions, especially from temporally sparse and noisy event streams given by event cameras in non-uniform motion.

In particular, we incorporate a more realistic event generation model that accounts for various intrinsic parameters (e.g. time-independent, asymmetric contrast threshold and refractory period) and non-idealities (e.g. pixel-to-pixel threshold variation). Furthermore, we introduce two complementary normalized reconstruction losses that are not only effectively invariant to the camera motion speed and threshold scale, but also minimally influenced by asymmetric thresholds. This allows for their effective generalization, as well as the regularization-free joint optimization of unknown contrast threshold and refractory period from poor initializations. The first loss serves as the primary loss for high-fidelity reconstruction, while the second acts as a smoothness constraint for better regularization of textureless regions. As both loss functions do not involve event accumulation, detailed and robust reconstruction from sparse and noisy events can be achieved. Our experiments on novel sequences, simulated using an improved variant of ESIM [36], and real sequences from TUM-VIE [17] verify the effectiveness of Robust *e*-NeRF. We publicly release our code, synthetic event dataset and improved ESIM.

2. Related Work

Event-based Scene Reconstruction. Successful reconstruction of geometry, and possibly appearance, from events has been demonstrated in notable works, such as [35, 10, 56, 37, 15, 57]. Nonetheless, reconstruction of scene appearance is largely limited to diffuse surfaces due to the adoption of Lambertian surface assumption via brightness constancy [37, 15]. Moreover, existing methods can mainly recover semi-dense geometry, in the form of discrete depth maps or point clouds, corresponding to edges (more precisely, locations with high perceived spatial intensity gradient) in the scene. This is due to the fact that events are primarily generated along edges under relative motion [9].

While [51] achieved dense diffuse reconstruction by applying the classic *Structure-from-Motion* (SfM) and *Multi-View Stereo* (MVS) pipelines on video frames reconstructed from events [38], its performance is intrinsically limited by the accuracy of the recovered frames, as similarly discussed in Sec. 1. Although limited to simple, object-level mesh or specialized parametric models, [42, 25, 53] also demonstrated dense diffuse reconstruction via *Analysis-by-Synthesis* or equivalently *Vision-as-Inverse-Graphics* [13]. In contrast, we aim to reconstruct *dense, continuous* scene geometry and *view-dependent* appearance, in the form of a NeRF, *directly* from the raw event stream.

Reconstructing Neural Radiance Fields. In general, visual reconstruction of neural scene representations, including NeRFs, is achieved via *Analysis-by-Synthesis*. Nonetheless, NeRF derivatives mainly focus on the reconstruction from dense [21, 4, 5] or sparse [26, 16] multi-view images, possibly with depth maps [2, 3] or point clouds [40, 6, 39].

The reconstruction of NeRFs from events was first proposed and investigated in Ev-NeRF [12], E-NeRF [18] and EventNeRF [41]. E-NeRF also explored on using a combination of events and images for NeRF reconstruction. However, these works suffer from various limitations, as outlined in Sec. 1. Moreover, EventNeRF also relies on the access to the analytic camera trajectory, in contrast to E-NeRF and our work which only require constant-rate camera poses. Furthermore, inconsistent sets of loss functions and hyper-parameters were also adopted across different scenes in E-NeRF. In addition, E-NeRF with normalized and no-event losses also requires the contrast threshold to be known as *a priori*, which is hard to achieve in practice.

3. Our Method

We first briefly introduce the *Neural Radiance Field* (NeRF) scene representation (Sec. 3.1). Next, we detail the event generation model (Sec. 3.2) and normalized training losses (Sec. 3.3) proposed to robustly reconstruct NeRFs from event cameras. Lastly, we describe a *Gamma Correction*-based approach to align the radiance levels of the synthesized views to a set of given reference views (Sec. 3.4).

3.1. Preliminaries: Neural Radiance Fields

Neural Radiance Field (NeRF) [24] represents a scene using a *Multi-Layer Perceptron* (MLP) $F_{\Theta} : (\mathbf{x}, \mathbf{d}) \mapsto (\mathbf{c}, \sigma)$ that maps 3D position $\mathbf{x} = (x, y, z)$ and 2D viewing direction $\mathbf{d} = (\theta, \phi)$ to its corresponding directional emitted radiance, or simply color, $\mathbf{c} = (r, g, b)$ and volume density σ . From this representation, the estimated incident radiance \hat{L} at a given pixel \mathbf{u} can be computed using the volume rendering equation with quadrature [46], as follows:

$$\hat{L}(\mathbf{u}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad (1)$$

$$T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right),$$

where σ_i and \mathbf{c}_i are the volume density and emitted radiance, respectively, of a sample \mathbf{x}_i along the back-projected ray through the pixel, which has direction \mathbf{d} from the camera center \mathbf{o} . The sample $\mathbf{x}_i = \mathbf{o} + s_i \mathbf{d}$ has a distance s_i from the camera center and a distance of $\delta_i = s_{i+1} - s_i$ between its adjacent sample \mathbf{x}_{i+1} .




Figure 2. Event generation model. An event e of polarity p is generated at timestamp t_{curr} when the difference in log-radiance $\log L$ at a pixel \mathbf{u} , measured with respect to a reference $\log L$ at timestamp t_{ref} , has the same sign as p and a magnitude that equals to the contrast threshold associated to polarity p , C_p . Red, downwards and blue, upwards arrows represent events of polarities -1 and $+1$, respectively, and each right-angled dashed line represents the measured change in $\log L$. After an event is generated, the pixel will be temporarily deactivated for an amount of time given by the refractory period τ , as shaded in the figure. Thus, t_{ref} is simply the sum of the previous event timestamp t_{prev} and τ .

3.2. Event Generation Model

Fig. 2 shows the illustration of the event generation model. An event camera responds to log-radiance changes in the scene and outputs an *Event Stream* \mathcal{E} , given by:

$$\mathcal{E} = \{ e \mid e = (\mathbf{u}, p, t_{prev}, t_{curr}) \}, \quad (2)$$

where e is an *Event* generated by pixel \mathbf{u} , with polarity $p \in \{-1, +1\}$, at timestamp t_{curr} . For convenience of discussion, we augment each event with the timestamp of the previous event that was generated by the same pixel, t_{prev} .

An event of polarity p is generated when the difference in log-radiance at a pixel, measured with respect to a reference log-radiance at timestamp t_{ref} , has the same sign as p and a magnitude that equals to the *Contrast Threshold* associated to polarity p , C_p [9]. In short, the condition is given by:

$$\Delta \log L := \log L(\mathbf{u}, t_{curr}) - \log L(\mathbf{u}, t_{ref}) = p C_p, \quad (3)$$

where L denotes the incident radiance at the given pixel and timestamp. For color event cameras, L corresponds to the radiance of the incident light after passing through the specific color filter in front of the pixel.

After an event is generated, the pixel will be deactivated for an amount of time specified by the *Refractory Period* τ . During this period of time, the pixel is invariant to any change in log-radiance and thus will not generate any new events. At the end of the refractory period, the pixel will be reactivated and the current log-radiance value at the pixel will be set as the new reference value, enabling the next event to be generated at this pixel [9, 19]. In essence:

$$t_{ref} = t_{prev} + \tau. \quad (4)$$

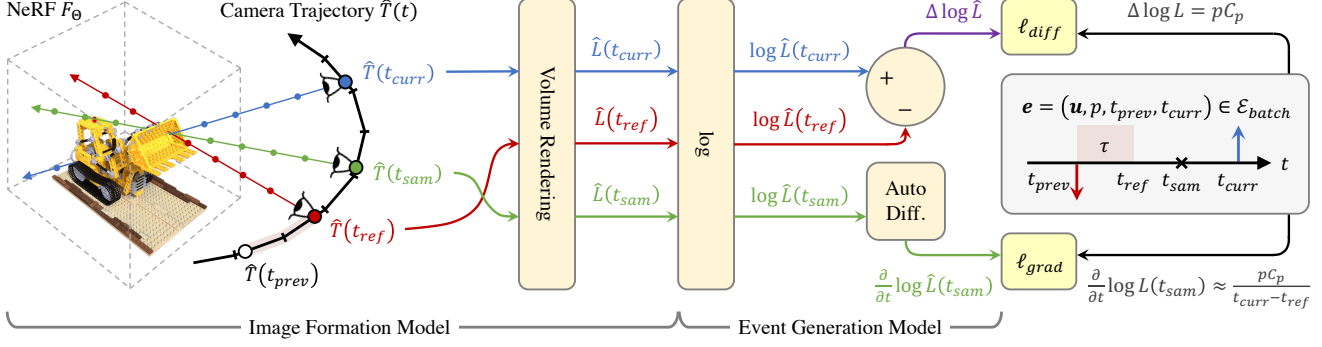


Figure 3. Overview of Robust e -NeRF training. For each event e in the batch \mathcal{E}_{batch} sampled randomly from the event stream, we first derive the reference timestamp t_{ref} (Eq. 4), given the refractory period τ , and sample a timestamp t_{sam} between t_{ref} and t_{curr} . Next, we interpolate the given constant-rate camera poses at t_{ref} , t_{sam} and t_{curr} using LERP for position and SLERP for orientation. Given these pose estimates \hat{T} , we then perform volume rendering on the back-projected rays from pixel \mathbf{u} with the NeRF F_{Θ} (Eq. 1). This is done to infer the predicted radiance \hat{L} , and thus log-radiance $\log \hat{L}$, of pixel \mathbf{u} at t_{ref} , t_{sam} and t_{curr} . For brevity, we denote $\hat{L}(t) = \hat{L}(\mathbf{u}, t)$. These $\log \hat{L}$ are ultimately used to derive the predicted log-radiance difference $\Delta \log \hat{L}$ and gradient $\frac{\partial}{\partial t} \log \hat{L}(t_{sam})$ for the computation of the reconstruction loss: threshold-normalized difference loss ℓ_{diff} (Eq. 6) and smoothness loss: target-normalized gradient loss ℓ_{grad} (Eq. 7), given the observed log-radiance difference $\Delta \log L$ (Eq. 3) and gradient $\frac{\partial}{\partial t} \log L(t_{sam})$ approximation from the event e , respectively.

The refractory period gives rise to temporal sparsity in the event stream. While this leads to partially observable log-radiance changes, it also limits the event generation rate of the camera [9, 19]. This is crucial to prevent the corruption of event timestamps due to *Address Event Representation* (AER) bus saturation and readout congestion, especially when the resolution and/or the speed of the camera are high [9, 1, 34, 8]. The bounded event rate also facilitates longer periods of recordings given a fixed memory budget. Furthermore, long refractory periods, as well as asymmetric contrast thresholds, contribute to a lower *Shot Noise* event rate [22], thereby improving *Signal-to-Noise Ratio* (SNR).

Event sensors also suffer from non-uniformity of pixel response, similar to standard image sensors. This is characterized by the pixel-to-pixel variation in the contrast threshold, which can be viewed as the event sensor analogue of *Fixed-Pattern Noise* (FPN). Studies suggest that the threshold variation can be modeled as a *Gaussian* distribution with mean C_p [9, 19, 33].

The contrast threshold is also effectively constant over time, as it is independent of temperature. The key assumption for temperature-independence is the absence of junction leakage [27], which holds true for modern event cameras [43]. Regardless of leakage, temperature-independence remains observed over a large temperature range [27].

3.3. Training

3.3.1 Assumptions

To reconstruct a NeRF from an event camera in motion, we assume that we are given the intrinsic camera matrix, lens distortion parameters and constant-rate camera poses of sufficiently high sampling rate for accurate interpolation

at arbitrary time instants. Similar to [18], we perform *Linear Interpolation* (LERP) on camera positions and *Spherical Linear Interpolation* (SLERP) on camera orientations. While our method does not rely on event camera-specific intrinsic parameters such as contrast threshold and refractory period to be known *a priori*, this information generally facilitates a more accurate reconstruction. As we will see in Sec. 3.3.2, it is also acceptable to provide only the (*positive-to-negative*) *contrast threshold ratio* C_{+1}/C_{-1} , which can be inferred from event camera bias settings.

3.3.2 Fundamental Limitations

As event cameras only provide (partial) observations of log-radiance *changes*, not *absolute* log-radiance, the predicted log-radiance $\log \hat{L}$, given by volume renderings from the reconstructed NeRF (Eq. 1), is only accurate up to an offset per color channel. Furthermore, there will be an additional scale ambiguity that is consistent across all channels, when only the contrast threshold ratio is known. This is similar to an image with unknown *black levels* and ISO. Nonetheless, these ambiguities can be easily dealt with or corrected for post-reconstruction, *e.g.* using a set of reference images of the same scene (Sec. 3.4), thereby not a cause for concern. We are also restricted to the reconstruction of a NeRF with single-channel directional emitted radiance, when a monochrome event camera is employed.

3.3.3 Loss Functions

An overview of the training pipeline is illustrated in Fig. 3. In line with the event generation model (Sec. 3.2), we propose two complementary normalized loss functions: *threshold-normalized difference loss* ℓ_{diff} and *target-*

normalized gradient loss ℓ_{grad} that directly and effectively generalize to various real-world conditions. The weighted sum of the two losses form the total training loss, which is optimized on a batch of events \mathcal{E}_{batch} sampled randomly from the raw, asynchronous event stream. Formally, the total training loss is given by:

$$\mathcal{L} = \frac{1}{|\mathcal{E}_{batch}|} \sum_{e \in \mathcal{E}_{batch}} \lambda_{diff} \ell_{diff}(e) + \lambda_{grad} \ell_{grad}(e), \quad (5)$$

where λ_{diff} and λ_{grad} are the respective loss weights.

We specifically refrained from optimizing on a reduced event stream, obtained by accumulating successive events at each pixel over time intervals, as done in [41, 12, 18]. This enables us to account for the refractory period and prevent unnecessary effective reduction in contrast sensitivity, which leads to lower reconstruction fidelity. Moreover, it can also be shown that event accumulation results in the effective amplification of threshold variation (proof in supplementary materials), thereby reduction in noise robustness.

Threshold-Normalized Difference Loss. This loss enforces the *mean contrast threshold* $\bar{C} = \frac{1}{2}(C_{-1} + C_{+1})$ normalized squared consistency between the observed log-radiance difference $\Delta \log L = pC_p$ from an event (Eq. 3) and the predicted log-radiance difference $\Delta \log \hat{L} := \log \hat{L}(\mathbf{u}, t_{curr}) - \log \hat{L}(\mathbf{u}, t_{ref})$, given by renders from the NeRF model (Eq. 1), as follows:

$$\ell_{diff}(e) = \left(\frac{\Delta \log \hat{L} - pC_p}{\bar{C}} \right)^2. \quad (6)$$

Note that when a color event camera is employed, \hat{L} refers to the single-channel rendered radiance, where the color channel is governed by the pixel color filter.

The loss serves as the primary reconstruction loss and can be effectively interpreted as a squared percentage error, especially under symmetric contrast thresholds. The normalization entails that the loss is invariant to the common scale of the positive and negative thresholds, as well as the predicted log-radiance, and only dependent on their ratio. Moreover, the normalization is optimal in the sense that the magnitude of the normalized target $|pC_p/\bar{C}| = C_p/\bar{C}$ is always centered at 1 regardless of the threshold ratio (proof in supplementary materials). The loss can therefore effectively generalize to arbitrary threshold values.

Unlike Ev-NeRF [12], these properties also enable the joint optimization of the unknown contrast threshold without additional regularization, as it does not suffer from any degeneracy. However, only the contrast threshold ratio can be recovered, thereby an additional scale ambiguity in the predicted log radiance, as mentioned in Sec. 3.3.2. In addition, our experiments also demonstrate the viability of jointly optimizing the refractory period τ via t_{ref} (Eq. 4).

Target-Normalized Gradient Loss. This loss is simply the *Absolute Percentage Error* (APE) of the predicted log-radiance temporal gradient derived using auto-differentiation $\frac{\partial}{\partial t} \log \hat{L}(\mathbf{u}, t)$, with respect to the finite difference approximation of the target log-radiance gradient $\frac{\partial}{\partial t} \log L(\mathbf{u}, t) \approx \frac{pC_p}{t_{curr} - t_{ref}}$, at a timestamp t_{sam} sampled between t_{ref} and t_{curr} , as follows:

$$\ell_{grad}(e) = \text{APE} \left(\frac{\partial}{\partial t} \log \hat{L}(\mathbf{u}, t_{sam}), \frac{pC_p}{t_{curr} - t_{ref}} \right) \quad (7)$$

where $\text{APE}(\hat{y}, y) = \left| \frac{\hat{y} - y}{y} \right|$. As the finite difference approximation error is minimum at the midpoint and maximum at the endpoints, we sample t_{sam} from a truncated normal distribution that is centered at the midpoint and has a standard deviation of $1/4$ the interval.

The loss acts as a smoothness constraint for log-radiance changes between t_{ref} and t_{curr} , which lack explicit regularization from ℓ_{diff} . This is particularly important for the effective reconstruction of textureless regions in the scene, where events associated have comparably longer intervals. In contrast to analogous regularization losses adopted in recent works [41, 12, 18], ℓ_{grad} is specifically invariant to the speed of motion, hence generalizable to arbitrary speed profiles. An unnormalized gradient loss would over-emphasize events generated under high-speed motion, as they have relatively larger target gradients. Furthermore, the loss is invariant to the common scale of the threshold and predicted log-radiance, similar to ℓ_{diff} . It is therefore also able to effectively generalize to arbitrary threshold values and facilitate joint optimization of unknown threshold.

3.4. Gamma Correction of Synthesized Views

As alluded in Sec. 3.3.2, the channel-consistent scale and per-channel offset ambiguity in the predicted log-radiance can be corrected for post-reconstruction, given a set of reference images of the same scene. To better account for the likely mismatch of *spectral sensitivity*, thereby *color balance*, between the event camera and the standard camera used to capture the reference images. Akin to [41], we further relax the channel-consistent scale to per-channel scales. This entails an *affine* correction of the predicted log-radiance, or equivalently a *gamma* correction of the predicted *linear* radiance, for each color channel as follows:

$$\log \hat{\mathbf{L}}_{corr} = \mathbf{a} \odot \log \hat{\mathbf{L}} + \mathbf{b}, \quad (8)$$

where \mathbf{a} and \mathbf{b} are the correction parameters, is necessary and sufficient for the reference alignment. Consequently, given the target log-radiance $\log \mathbf{L}$ from the reference images, the optimal correction parameters can be simply derived via *ordinary least squares*.

4. Experiments

We adopt *Novel View Synthesis* (NVS) as the standard task to verify that our method, Robust e -NeRF can indeed directly and robustly reconstruct NeRFs from moving event cameras under various real-world conditions, particularly from sparse and noisy events generated under non-uniform motion. The NVS benchmark experiments are conducted on both synthetic (Sec. 4.1) and real sequences (Sec. 4.2). In addition, we perform ablation studies (Sec. 4.3) to investigate the significance of various components in our method.

Metrics. We adopt a consistent set of metrics to assess the performance of a method in all experiments. In particular, we employ the three following standard metrics: *Peak Signal-to-Noise Ratio* (PSNR), *Structural Similarity Index Measure* (SSIM) [50] and AlexNet-based *Learned Perceptual Image Patch Similarity* (LPIPS) [55] to quantify the similarity between the gamma-corrected (Sec. 3.4) synthesized novel views and given target novel views.

Baselines. We benchmark our method, Robust e -NeRF against a recent work, Ev-NeRF [12] and a naïve baseline, E2VID + NeRF, given by cascading the seminal events-to-video reconstruction method, E2VID [38] to NeRF. For Ev-NeRF, events are accumulated over non-overlapping time intervals of $1/24 s \approx 41.67 ms$, as suggested in their paper. To facilitate a fair comparison, all methods, including ours, have been (re-)implemented to adopt a common NeRF backbone. While we do not explicitly compare against E-NeRF [18] and EventNeRF [41], our experiment results should still provide an accurate indication on their performance relative to ours since they share a lot of similarities with Ev-NeRF as detailed in Sec. 1.

Datasets. The synthetic experiments are performed on a novel set of sequences simulated on the “Realistic Synthetic 360°” scenes, which were proposed in the seminal NeRF paper [24]. These scenes contain a wide variety of photo-realistic objects with complicated structure and non-Lambertian effects, thereby effective for NVS evaluation. The new synthetic event dataset allows for a retrospective comparison between event-based and image-based NeRF reconstruction methods, as the sequences were simulated under highly similar conditions, unlike in [41, 12].

The events are generated from a virtual event camera moving in a hemi-/spherical spiral motion about the object at the origin, as shown in Fig. 1(b, d), using an improved variant of ESIM [36], which is an efficient and realistic event simulator. Specifically, we added support to time-independent pixel-to-pixel threshold variation and Blender, circumvented singularities in the trajectory orientation interpolation, as well as improved event simulation accuracy, especially with non-zero refractory periods. On the contrary, the real experiments are performed on the

`mocap-1d-trans` and `mocap-desk2` sequences of the TUM-VIE dataset [17], which are mainly forward-facing captures of a desk with various objects on top, under linear and spiral camera motion, respectively. These real sequences were chosen for their relative suitability for NVS and their adoption of a modern, high-resolution event sensor — Prophesee Gen 4.

4.1. Synthetic Experiments

The synthetic experiments form the core of the benchmark, as they allow for realistic *controlled* experiments under various real-world conditions with *absolute* ground truth, which are otherwise infeasible using real sequences.

All sequences are simulated with a symmetric contrast threshold of 0.25 (*i.e.* $C_{-1} = C_{+1} = 0.25$), which is the approximate nominal threshold for event sensors such as the Prophesee Gen 3.1, Gen 4.1 and Sony IMX636 sensor, hence providing a known threshold ratio of 1 (*i.e.* $C_{+1}/C_{-1} = 1$). The event camera trajectory is also sampled at a high rate of 1 *kHz* to minimize the influence of pose errors on the performance assessment. Furthermore, the virtual event camera revolves the object 4 times, with uniform 1 revolution per second speed about the object vertical axis by default, similar to [41]. Unless otherwise stated, a sequence is also simulated with zero pixel-to-pixel threshold standard deviation and refractory period (*i.e.* $\sigma_{C_p} = 0, \tau = 0$). By default, Ev-NeRF and our method are trained with a constant symmetric contrast threshold, given by the prior knowledge of time-independent threshold ratio.

Effect of Speed Profile. To investigate the influence of camera speed, we evaluate all works on 4 sets of sequences simulated with different speed profiles, where each set contains a sequence for each of the 7 scenes. Particularly, the 1st set is simulated with the default settings, thus providing a reference (uniform azimuth) speed of motion (*i.e.* $v = 1\times$). We oscillate the speed of motion between $1/8\times$ and $8\times$ the original speed at a frequency of 1Hz (*i.e.* $v = v_b \sin 2\pi ft$, where $v_b = 8\times$ and $f = 1Hz$) for the 2nd set, and scale the speed of motion to $1/8\times$ and $8\times$ of the original speed (*i.e.* $v = 1/8\times, 8\times$) for the 3rd and 4th sets, respectively. Fig. 1(b) illustrates the uniform motion in the 1st, 3rd and 4th sets, whereas Fig. 1(d) illustrates the non-uniform, oscillating motion in the 2nd set.

The quantitative results reported in Tab. 1 underscores the significance of speed invariance in enabling the effective generalization to arbitrary speed profiles, as it can be observed that the performance of Ev-NeRF deteriorates as the speed deviates from the optimal at $v = 1\times$. Our method also outperforms all baselines, including Ev-NeRF, at the default setting (*i.e.* $v = 1\times$), which is optimal for all. Qualitative results shown in Fig. 4 suggests an overall improvement in synthesis quality, including high-frequency details.

Method	$v = 1\times$			$v_b = 8\times$			$v = \frac{1}{8}\times$			$v = 8\times$		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
E2VID + NeRF	18.92	0.832	0.316	18.92	0.832	0.316	18.92	0.832	0.316	18.92	0.832	0.316
Ev-NeRF	27.72	0.935	0.087	26.25	0.926	0.102	19.79	0.792	0.326	20.83	0.862	0.198
Robust e -NeRF	28.19	0.945	0.057	28.19	0.945	0.057	28.19	0.945	0.057	28.19	0.945	0.057

Table 1. Effect of speed profile. v denotes the speed of motion relative to the default hemi-/spherical spiral motion with uniform azimuth speed, whereas v_b denotes the oscillation factor of the relative speed of motion (*i.e.* $v = v_b \sin 2\pi ft$, $f = 1Hz$).

Method	Opt. C_p	$\sigma_{C_p} = 0.00$			$\sigma_{C_p} = 0.03$			$\sigma_{C_p} = 0.06$		
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
E2VID + NeRF	–	18.92	0.832	0.316	18.68	0.827	0.330	18.03	0.808	0.363
Ev-NeRF	×	27.72	0.935	0.087	24.42	0.895	0.155	8.07	0.841	0.260
	✓	27.43	0.911	0.123	23.66	0.826	0.261	15.43	0.708	0.441
Robust e -NeRF	×	28.19	0.945	0.057	28.14	0.946	0.058	28.23	0.947	0.057
	✓	28.17	0.946	0.051	27.91	0.946	0.054	28.19	0.948	0.049

Table 2. Effect of pixel-to-pixel threshold variation σ_{C_p} . “Opt. C_p ” refers to jointly optimizing thresholds C_p with NeRF parameters Θ .

Method	Opt. C_p	Opt. τ	$\tau = 0ms$			$\tau = 8ms$			$\tau = 25ms$		
			PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
E2VID + NeRF	–	–	18.92	0.832	0.316	14.87	0.797	0.427	14.15	0.791	0.467
Ev-NeRF	×	–	27.72	0.935	0.087	13.17	0.707	0.559	12.75	0.759	0.528
	✓	–	27.43	0.911	0.123	13.56	0.716	0.528	13.75	0.717	0.569
Robust e -NeRF	×	×	28.19	0.945	0.057	26.30	0.934	0.066	25.51	0.929	0.072
	×	✓	28.18	0.945	0.052	23.43	0.910	0.090	22.48	0.895	0.105

Table 3. Effect of refractory period τ . “Opt. τ ” refers to jointly optimizing refractory period τ with NeRF parameters Θ .

Method	Opt. C_p	Opt. τ	$v_b = 1\times, \sigma_{C_p} = 0.00, \tau = 0ms$			$v_b = 4\times, \sigma_{C_p} = 0.03, \tau = 8ms$			$v_b = 8\times, \sigma_{C_p} = 0.06, \tau = 25ms$		
			PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
E2VID + NeRF	–	–	18.92	0.832	0.316	14.98	0.796	0.433	14.07	0.801	0.448
Ev-NeRF	×	–	27.72	0.935	0.087	12.33	0.742	0.521	12.05	0.807	0.425
	✓	–	27.43	0.911	0.123	13.06	0.732	0.539	12.27	0.772	0.539
Robust e -NeRF	×	×	28.19	0.945	0.057	24.10	0.913	0.086	23.51	0.900	0.110
	✓	✓	28.19	0.946	0.051	20.42	0.875	0.126	18.83	0.836	0.197

Table 4. Collective effect of speed profile, threshold variation and refractory period.

τ	ℓ_{grad}	$v_b = 1\times, \sigma_{C_p} = 0.00, \tau = 0ms$			$v_b = 4\times, \sigma_{C_p} = 0.03, \tau = 8ms$			$v_b = 8\times, \sigma_{C_p} = 0.06, \tau = 25ms$		
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
×	✓	28.19	0.945	0.057	12.77	0.799	0.372	12.41	0.798	0.412
✓	×	27.96	0.943	0.063	23.15	0.899	0.113	22.21	0.879	0.153
✓	✓	28.19	0.945	0.057	24.10	0.913	0.086	23.51	0.900	0.110

Table 5. Ablation studies on refractory period τ modeling and target-normalized gradient loss ℓ_{grad} .

Effect of Pixel-to-Pixel Threshold Variation. To evaluate the robustness of our method to noise, in the form of threshold variation, we benchmark all works on three sets

of sequences. Furthermore, we also benchmark Ev-NeRF and our method with jointly optimized contrast thresholds, which are poorly initialized with $C_{+1}/C_{-1} = 10$ (more




Figure 4. Synthesized novel views in the synthetic and real experiments. Results shown correspond to that with jointly optimized contrast thresholds and possibly refractory period, if applicable for the method in the particular experiment.

precisely, we set $C_{-1} = 0.25$ and $C_{+1} = 2.5$), to assess the robustness of the joint optimization. The three sets of sequences were simulated with pixel-to-pixel threshold standard deviations of 0, 0.03 and 0.06 (*i.e.* $\sigma_{C_p} = 0.00, 0.03, 0.06$), which are equivalent to 0%, 12% and 24% of the contrast threshold, respectively. $\sigma_{C_p} = 0.03$ is approximately the nominal value for the Prophesee Gen 3.1, Gen 4 and Sony IMX636 sensor, and $\sigma_{C_p} = 0.06$ is the maximum value for the Sony IMX636 sensor.

The quantitative results given in Tab. 2 clearly demonstrates the robustness of our method to threshold variation, as our performance is essentially unaffected by the degree of noise. In contrast, the baselines are visibly susceptible to threshold variation, as their performance declines with its severity, especially Ev-NeRF. It is interesting to note that even the naïve baseline outperforms Ev-NeRF at $\sigma_{C_p} = 0.06$. These conclusions are also supported qualitatively in

Fig. 4. The quantitative results also illustrate the effectiveness and robustness of our threshold joint optimization, as the performance of our method with and without jointly optimized thresholds are virtually the same. In contrast, it can be observed that the joint optimization of time-varying contrast thresholds in Ev-NeRF leads to a slight decrease in performance on the default setting (*i.e.* $\sigma_{C_p} = 0.00$).

Effect of Refractory Period. To study the effect of temporal sparsity in the event stream due to the refractory period, we benchmark all methods on three sets of sequences simulated with refractory periods of 0, 8 and 25ms (*i.e.* $\tau = 0, 8, 25ms$). Similar to the previous experiment, we additionally evaluate our method with jointly optimized refractory period, which is poorly initialized with half the maximum possible value, given by the minimum time interval between successive events at a pixel in the sequence.

Furthermore, we also benchmark Ev-NeRF with jointly optimized thresholds to validate the importance of accounting for the refractory period. $\tau = 8ms$ is just slightly less than mean event interval in the set of sequences with $\tau = 0ms$. Let \mathcal{E}_τ be the event stream with refractory period τ , the degree of sparsity due to τ can be quantified by the mean $|\mathcal{E}_{\tau=0ms}|/|\mathcal{E}_\tau|$ across all sequences in the set, which translates to $5.84\times$ and $11.37\times$ for $\tau = 8$ and $25ms$, respectively.

The quantitative results presented in Tab. 3 undoubtedly verifies the importance of refractory period modeling in enabling the robust reconstruction from temporally sparse event streams, as our method significantly outperforms all baselines across all values of τ . Moreover, the attempt by Ev-NeRF to compensate τ with jointly optimized thresholds has also clearly failed, albeit contributing to slightly improved performance. It is also worth noting that the naïve baseline achieves better reconstructions than Ev-NeRF under non-zero τ . Our perceivable drop in performance, as τ increases, may be attributed to its associated decrease in observability. Our results also demonstrate the feasibility of jointly optimizing τ , although less effective than threshold joint optimization, which reflects the complexities involved.

Collective Effect. To assess the performance of all methods under the collective effect of threshold variation, refractory period and camera speed, which resembles real-world operating conditions, we benchmark all works on three sets of sequences with different levels of difficulty: *easy* ($\sigma_{C_p} = 0.00, \tau = 0ms, v_b = 1\times$), *medium* ($\sigma_{C_p} = 0.03, \tau = 8ms, v_b = 4\times$) and *hard* ($\sigma_{C_p} = 0.06, \tau = 25ms, v_b = 8\times$). In addition, we also benchmark Ev-NeRF and our method with jointly optimized intrinsic parameters, similar to previous experiments. The medium and hard sets are $5.31\times$ and $8.27\times$ sparser than the easy set, respectively. This experiment is particularly challenging as a $K\times$ speed-up effectively scales τ by $K\times$ in the context of the original sequence. Moreover, methods that involve event accumulation will also experience a $K\times$ effective scaling in σ_{C_p} .

The quantitative results provided in Tab. 4 generally reflect that of the previous experiment, thus the same conclusions can be drawn. This should not be surprising, given the proven importance of accounting for τ . Our decline in performance on the medium and hard settings, when compared to $\tau = 8$ and $25ms$, may be attributed to the complications arising from the interaction between σ_{C_p} , τ and v_b . Qualitative results given in Fig. 4 illustrates the strength and robustness of our method under challenging conditions.

4.2. Real Experiments

The real experiments mainly serve as a qualitative benchmark, as the dataset (and other similar ones) is not specifically suited for the task of NVS, thereby unable to provide an accurate quantification of performance (refer to supplementary materials for a detailed justification). For the real

experiments, Ev-NeRF and our method are trained with jointly optimized intrinsic parameters, where the contrast threshold is initialized with $C_{+1}/C_{-1} = 1$ (more precisely, we set $C_{-1} = C_{+1} = 0.25$) this time. While the two sequences mainly involve an insignificant refractory period relative to their primarily non-uniform speed of motion, qualitative results shown in Fig. 4 still demonstrate our superior performance in recovering fine details while appropriately smoothing uniform regions, in line with the synthetic experiments. Exposure of the scene is also visibly more consistent and more accurately reconstructed. Visible artifacts near the borders of all synthesized views are due to the comparably narrower *field-of-view* of the event camera.

4.3. Ablation Studies

The ablation studies are conducted in the same setting as the synthetic experiment on the collective effect, without jointly optimized intrinsic parameters. The quantitative results reported in Tab. 5 further validates the immense importance of accounting for τ , even when event accumulation is not involved. Although the synthetic sequences mainly involves scenes with abundant texture, the results also verify the effectiveness of the target-normalized gradient loss in regularizing textureless regions, particularly under more challenging conditions (*i.e.* medium and hard settings).

5. Conclusion

In this paper, we introduce Robust *e*-NeRF, a novel approach to directly and robustly reconstruct NeRFs from moving event cameras under various real-world conditions, particularly from sparse and noisy events generated under non-uniform motion. Our method consists of two key components: a realistic event generation model that accounts for various event camera-specific intrinsic parameters and non-idealities, as well as a complementary pair of normalized reconstruction losses that can effectively generalize to arbitrary speed profiles and intrinsic parameter values without such prior knowledge. The proposed *Analysis-by-Synthesis* approach for event-based reconstruction naturally extends to other scene representations, such as 3D Gaussians [14]. In spite of its achievements, Robust *e*-NeRF still relies on the assumption of a static scene and given constant-rate camera poses, which we leave for future work. We also see the modeling of other spatial/temporal noise and artifacts of an event sensor as a promising future direction for more accurate reconstructions.

Acknowledgements. This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-RP-2021-024), and the Tier 2 grant MOE-T2EP20120-0011 from the Singapore Ministry of Education.

References

- [1] Understanding the performance of neuromorphic event-based vision sensors. Technical report, iniVation AG, 2020. [4](#)
- [2] Benjamin Attal, Eliot Laidlaw, Aaron Gokaslan, Changil Kim, Christian Richardt, James Tompkin, and Matthew O’Toole. TöRF: Time-of-Flight Radiance Fields for Dynamic Scene View Synthesis. In *Advances in Neural Information Processing Systems*, 2021. [3](#)
- [3] Dejan Azinović, Ricardo Martin-Brualla, Dan B Goldman, Matthias Nießner, and Justus Thies. Neural RGB-D Surface Reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [3](#)
- [4] J T Barron, B Mildenhall, M Tancik, P Hedman, R Martin-Brualla, and P P Srinivasan. Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [3](#)
- [5] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [3](#)
- [6] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised {NeRF}: Fewer Views and Faster Training for Free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022. [3](#)
- [7] Davide Falanga, Kevin Kleber, and Davide Scaramuzza. Dynamic obstacle avoidance for quadrotors with event cameras. *Science Robotics*, 2020. [2](#)
- [8] Thomas Finateau, Atsumi Niwa, Daniel Matolin, Koya Tsuchimoto, Andrea Mascheroni, Etienne Reynaud, Poooria Mostafalu, Frederick Brady, Ludovic Chotard, Florian LeGoff, Hirotsugu Takahashi, Hayato Wakabayashi, Yusuke Oike, and Christoph Posch. 5.10 A 1280×720 Back-Illuminated Stacked Temporal Contrast Event-Based Vision Sensor with 4.86µm Pixels, 1.066GEPS Readout, Programmable Event-Rate Controller and Compressive Data-Formatting Pipeline. In *2020 IEEE International Solid-State Circuits Conference - (ISSCC)*, 2020. [4](#)
- [9] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J. Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. [1](#), [2](#), [3](#), [4](#)
- [10] Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. A Unifying Contrast Maximization Framework for Event Cameras, with Applications to Motion, Depth, and Optical Flow Estimation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018. [2](#)
- [11] Javier Hidalgo-Carrió, Guillermo Gallego, and Davide Scaramuzza. Event-Aided Direct Sparse Odometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [2](#)
- [12] Inwoo Hwang, Junho Kim, and Young Min Kim. Ev-nerf: Event based neural radiance field. In *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2023. [2](#), [3](#), [5](#), [6](#)
- [13] Hiroharu Kato, Deniz Beker, Mihai Morariu, Takahiro Ando, Toru Matsuoka, Wadim Kehl, and Adrien Gaidon. Differentiable Rendering: A Survey, 2020. arXiv:2006.12057 [cs]. [2](#)
- [14] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.*, 2023. [9](#)
- [15] Hanme Kim, Stefan Leutenegger, and Andrew Davison. Real-time 3d reconstruction and 6-dof tracking with an event camera. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VI*, 2016. [2](#)
- [16] Mijeong Kim, Seonguk Seo, and Bohyung Han. InfoNeRF: Ray Entropy Minimization for Few-Shot Neural Volume Rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [3](#)
- [17] Simon Klenk, Jason Chui, Nikolaus Demmel, and Daniel Cremers. Tum-vie: The tum stereo visual-inertial event dataset. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021. [2](#), [6](#)
- [18] Simon Klenk, Lukas Koestler, Davide Scaramuzza, and Daniel Cremers. E-nerf: Neural radiance fields from a moving event camera. *IEEE Robotics and Automation Letters*, 2023. [2](#), [3](#), [4](#), [5](#), [6](#)
- [19] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. A 128 × 128 120 dB 15 µs latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits*, 2008. [3](#), [4](#)
- [20] Weng Fei Low and Gim Hee Lee. Minimal Neural Atlas: Parameterizing Complex Surfaces with Minimal Charts and Distortion. In *European Conference on Computer Vision (ECCV)*, 2022. [2](#)
- [21] 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*, 2021. [3](#)
- [22] Brian McReynolds, Rui Graca, and Tobi Delbruck. Exploiting Alternating DVS Shot Noise Event Pair Statistics to Reduce Background Activity. In *2023 International Image Sensor Workshop (IISW)*, 2023. [4](#)
- [23] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [2](#)
- [24] 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*, 2020. [2](#), [3](#), [6](#)

- [25] Jalees Nehvi, Vladislav Golyanik, Franziska Mueller, Hans-Peter Seidel, Mohamed Elgharib, and Christian Theobalt. Differentiable event stream simulator for non-rigid 3d tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 2
- [26] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi S M Sajjadi, Andreas Geiger, and Noha Radwan. RegNeRF: Regularizing Neural Radiance Fields for View Synthesis from Sparse Inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [27] Yuji Nozaki and Tobi Delbruck. Temperature and Parasitic Photocurrent Effects in Dynamic Vision Sensors. *IEEE Transactions on Electron Devices*, 2017. 4
- [28] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction, 2021. 2
- [29] F. Paredes-Valles and G. E. de Croon. Back to event basics: Self-supervised learning of image reconstruction for event cameras via photometric constancy. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [30] 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 Trans. Graph.*, 2021. 2
- [31] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [32] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [33] Christoph Posch and Daniel Matolin. Sensitivity and uniformity of a $0.18\mu\text{m}$ CMOS temporal contrast pixel array. In *2011 IEEE International Symposium of Circuits and Systems (ISCAS)*, 2011. 4
- [34] Christoph Posch, Teresa Serrano-Gotarredona, Bernabe Linares-Barranco, and Tobi Delbruck. Retinomorph Event-Based Vision Sensors: Bioinspired Cameras With Spiking Output. *Proceedings of the IEEE*, 2014. 4
- [35] Henri Rebecq, Guillermo Gallego, Elias Mueggler, and Davide Scaramuzza. EMVS: Event-Based Multi-View Stereo—3D Reconstruction with an Event Camera in Real-Time. *International Journal of Computer Vision*, 2018. 2
- [36] Henri Rebecq, Daniel Gehrig, and Davide Scaramuzza. ESIM: an Open Event Camera Simulator. In *Proceedings of The 2nd Conference on Robot Learning, PMLR*, 2018. 2, 6
- [37] 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. *IEEE Robotics and Automation Letters*, 2016. 2
- [38] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. High speed and high dynamic range video with an event camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 2, 6
- [39] Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Tom Funkhouser, and Vittorio Ferrari. Urban Radiance Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [40] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense Depth Priors for Neural Radiance Fields from Sparse Input Views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [41] Viktor Rudnev, Mohamed Elgharib, Christian Theobalt, and Vladislav Golyanik. Eventnerf: Neural radiance fields from a single colour event camera, 2022. arXiv:2206.11896 [cs]. 2, 3, 5, 6
- [42] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [43] Bongki Son, Yunjae Suh, Sungho Kim, Heejae Jung, Jun-Seok Kim, Changwoo Shin, Keunju Park, Kyoobin Lee, Jinman Park, Jooyeon Woo, Yohan Roh, Hyunku Lee, Yibing Wang, Ilia Ovsianikov, and Hyunsurk Ryu. 4.1 A 640×480 dynamic vision sensor with a $9\mu\text{m}$ pixel and 300Meps address-event representation. In *2017 IEEE International Solid-State Circuits Conference (ISSCC)*, 2017. 4
- [44] Timo Stoffregen, Cedric Scheerlinck, Davide Scaramuzza, Tom Drummond, Nick Barnes, Lindsay Kleeman, and Robert Mahony. Reducing the Sim-to-Real Gap for Event Cameras. In *Computer Vision – ECCV 2020*, 2020. 2
- [45] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [46] Andrea Tagliasacchi and Ben Mildenhall. Volume Rendering Digest (for NeRF), 2022. arXiv:2209.02417 [cs]. 3
- [47] Stepan Tulyakov, Daniel Gehrig, Stamatios Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza. Time lens: Event-based video frame interpolation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [48] Antoni Rosinol Vidal, Henri Rebecq, Timo Horstschaefer, and Davide Scaramuzza. Ultimate slam? combining events, images, and imu for robust visual slam in hdr and high-speed scenarios. *IEEE Robotics and Automation Letters*, 2018. 2
- [49] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *Advances in Neural Information Processing Systems*, 2021. 2
- [50] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004. 6
- [51] Kun Xiao, Guohui Wang, Yi Chen, Jinghong Nan, and Yongfeng Xie. Event-based dense reconstruction pipeline.

- In *2022 6th International Conference on Robotics and Automation Sciences (ICRAS)*, 2022. 2
- [52] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural Fields in Visual Computing and Beyond. *Computer Graphics Forum*, 2022. 2
- [53] Yuxuan Xue, Haolong Li, Stefan Leutenegger, and Joerg Stueckler. Event-based non-rigid reconstruction from contours. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*, 2022. 2
- [54] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Advances in Neural Information Processing Systems*, 2021. 2
- [55] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [56] Yi Zhou, Guillermo Gallego, Henri Rebecq, Laurent Kneip, Hongdong Li, and Davide Scaramuzza. Semi-dense 3d reconstruction with a stereo event camera. In *Computer Vision – ECCV 2018: 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part I*, 2018. 2
- [57] Yi Zhou, Guillermo Gallego, and Shaojie Shen. Event-based stereo visual odometry. *IEEE Transactions on Robotics*, 2021. 2