# Deblur *e*-NeRF: NeRF from Motion-Blurred Events under High-speed or Low-light Conditions

Weng Fei Low<sup>®</sup> and Gim Hee Lee<sup>®</sup>

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/deblur-e-nerf

Abstract. The distinctive design philosophy of event cameras makes them ideal for high-speed, high dynamic range & low-light environments, where standard cameras underperform. However, event cameras also suffer from motion blur, especially under these challenging conditions, contrary to what most think. This is due to the limited bandwidth of the event sensor pixel, which is mostly proportional to the light intensity. Thus, to ensure event cameras can truly excel in such conditions where it has an edge over standard cameras, event motion blur must be accounted for in downstream tasks, especially reconstruction. However, no prior work on reconstructing Neural Radiance Fields (NeRFs) from events, nor event simulators, have considered the full effects of event motion blur. To this end, we propose, Deblur *e*-NeRF, a novel method to directly and effectively reconstruct blur-minimal NeRFs from motion-blurred events. generated under high-speed or low-light conditions. The core component of this work is a physically-accurate pixel bandwidth model that accounts for event motion blur. We also introduce a threshold-normalized total variation loss to better regularize large textureless patches. Experiments on real & novel realistically simulated sequences verify our effectiveness. Our code, event simulator and synthetic event dataset are open-sourced.

Keywords: Neural Radiance Field  $\cdot$  Motion blur  $\cdot$  Event camera

## 1 Introduction

Event cameras offer a complementary approach to visual sensing, commonly achieved with frame-based cameras. Instead of capturing intensity images at a fixed rate, event cameras asynchronously detect changes in log-intensity per pixel and output a stream of *events*, each encoding the time instant, pixel location and direction of change. Such a stark contrast in design philosophy enable event cameras to offer many attractive properties over standard cameras, *e.g.* high dynamic range, high temporal resolution, low latency and low power [7].

These desirable properties make event cameras particularly ideal for applications that involve high-speed motion, *High Dynamic Range* (HDR)/low-light scenes and/or a strict power budget, such as in robotics, augmented reality,



Fig. 1: Event motion blur from a white bar moving on a black background (From [12])

surveillance and mobile imaging. Since these are exactly the operating conditions where standard cameras underperform, event cameras meaningfully complements standard cameras. This is clearly demonstrated in the recent success of image deblurring [14, 25, 46, 53], attributed to the addition of an event camera.

Nonetheless, event cameras also suffer from motion blur [12,26,49], especially under high speed or low light, albeit much less severe than standard cameras. This is contrary to what most think, as it has not been widely documented and discussed in the computer vision community. In general, motion blur of events are manifested as a time-varying latency on the event generation process. In severe cases, a significant "loss or introduction of events" may also occur, especially the former. This leads to artifacts such as event trails and blurring of edges, when visualizing events in 2D as their image-plane projection, as shown in Fig. 1. Event motion blur can be attributed to the limited bandwidth of the event sensor pixel, which is mostly proportional to the incident light intensity [3, 5, 8, 9, 12, 23, 30]. It bounds the minimum event detection latency, maximum frequency of change detectable and hence maximum event generation rate.

Therefore, to ensure event cameras can truly excel under conditions of highspeed or low-light, where it has an edge over standard cameras, it is crucial to account for event motion blur in downstream tasks, especially reconstruction. While recent works on reconstructing *Neural Radiance Fields* (NeRFs) [31] from events [13, 17, 27, 39], and possibly images [1, 17, 29, 36], have shown impressive results, none of them accounted for event motion blur, which limits their performance. Moreover, none of the existing event simulators [12, 15, 37] model the full non-linear behaviour of the pixel bandwidth under arbitrary lighting conditions. They at most model the 1<sup>st</sup>-order behavior exhibited under extreme low light.

**Contributions.** We propose Deblur *e*-NeRF, a novel method to directly and effectively reconstruct blur-minimal NeRFs from motion-blurred events, generated under high-speed motion or low-light conditions.

Specifically, we introduce a physically-accurate pixel bandwidth model to account for event motion blur under arbitrary speed & lighting conditions. We also present a discrete-time variant of the model, a numerical solution to its transient response & an importance sampling strategy, to enable its computational implementation. We incorporate them as part of the event generation model to reconstruct blur-minimal NeRFs via *Analysis-by-Synthesis* [16], which also supports the joint optimization of unknown pixel bandwidth model parameters. We also introduce a novel threshold-normalized total variation loss to better regularize large textureless patches in the scene. Ambiguities in the reconstruction are resolved by performing the proposed translated-gamma correction, which takes the pixel bandwidth model into consideration. Experiments on new event sequences, simulated with an improved ESIM [27,37] using our pixel bandwidth model, & real sequences from EDS [11], clearly validate the effectiveness of Deblur *e*-NeRF. Our code, event simulator & synthetic event dataset are open-sourced.

# 2 Related Work

**Image Motion Deblurring.** Image motion blur can be simply modeled as an average of incident light intensity over the exposure time of the image [35], which is a unity-gain *Linear Time-Invariant* (LTI) *Low-Pass Filter* (LPF). Thus, the severity of image motion blur is invariant to lighting, unlike event motion blur.

This model may be used in an *Analysis-by-Synthesis* framework [16] to deblur images [33]. However, a simplified model of spatially filtering images with a time-varying blur kernel [35] is more commonly used for this task [2,18,19,34,40,47,48]. With the rise of deep learning, state-of-the-art image/video deblurring methods [41,42,44,50,51] are generally deep image-to-image translation networks.

Despite the rich literature on image motion deblurring, there are no known methods for effective event motion deblurring. While concurrent work [49] proposed to correct event timestamps for motion blur-induced latency (Sec. 1), the method cannot handle any blur-induced "loss or introduction of events" and does not generalize to different cameras and bias settings. Similar generalization issues also plague concurrent work [26] on night-time events-to-video reconstruction.

**NeRF from Motion-Blurred Images.** A naïve way to reconstruct a blurminimal *Neural Radiance Field* (NeRF) [31] from motion-blurred images is to first deblur them using an existing method. Recent works have shown superior performance by integrating either the full [21, 43] or simplified [20, 28] image motion blur model into an Analysis-by-Synthesis framework, where both NeRF and blur model parameters are jointly optimized, similar to our work. However, there are no known works on NeRF reconstruction from motion-blurred events.

# 3 Preliminaries: Robust *e*-NeRF

Robust e-NeRF [27] is the state-of-the-art for reconstructing NeRFs with event cameras, particularly from temporally sparse and noisy events generated under non-uniform motion. The method consists of 2 key components: a realistic event generation model and a pair of normalized reconstruction losses. After training, NeRF renders are *gamma corrected* to resolve ambiguities in the reconstruction.

**Event Generation Model.** An *Event*, denoted as  $e = (u, p, t_{prev}, t_{curr})$ , with polarity  $p \in \{-1, +1\}$  is generated at timestamp  $t_{curr}$  when the change in incident log-radiance log L at a pixel u, measured relative to a reference value at

timestamp  $t_{ref}$ , shares the same sign as p and possesses a magnitude given by the *Contrast Threshold* associated to polarity p,  $C_p$ . Following the generation of an event, the pixel will be momentarily deactivated for an amount of time determined by the *Refractory Period*  $\tau$  and then reset at the end. This event generation model, as illustrated in Fig. 2, can be succinctly described by:

$$\Delta \log L := \log L(\boldsymbol{u}, t_{curr}) - \log L(\boldsymbol{u}. t_{ref}) = pC_p \text{, where } t_{ref} = t_{prev} + \tau \text{.} \quad (1)$$

**Training.** To reconstruct a NeRF from an *Event Stream*  $\mathcal{E} = \{e\}$ , provided by a calibrated event camera with known trajectory, a batch of events  $\mathcal{E}_{batch}$  is sampled randomly from  $\mathcal{E}$  for optimization of the following total training loss:

$$\mathcal{L} = \frac{1}{|\mathcal{E}_{batch}|} \sum_{\boldsymbol{e} \in \mathcal{E}_{batch}} \lambda_{diff} \ell_{diff}(\boldsymbol{e}) + \lambda_{grad} \ell_{grad}(\boldsymbol{e}) .$$
(2)

The threshold-normalized difference loss  $\ell_{diff}$  with weight  $\lambda_{diff}$  acts as the main reconstruction loss. It enforces the mean contrast threshold  $\tilde{C} = \frac{1}{2}(C_{-1} + C_{+1})$ normalized squared consistency between the predicted log-radiance difference  $\Delta \log \hat{L} := \log \hat{L}(\boldsymbol{u}, t_{curr}) - \log \hat{L}(\boldsymbol{u}, t_{ref})$ , given by NeRF renders, and the observed log-radiance difference  $\Delta \log L = pC_p$  from an event (Eq. (1)), as follows:

$$\ell_{diff}(\boldsymbol{e}) = \left(\frac{\Delta \log \hat{L} - pC_p}{\bar{C}}\right)^2 \,. \tag{3}$$

The target-normalized gradient loss  $\ell_{grad}$  with weight  $\lambda_{grad}$  acts as a smoothness constraint for regularization of textureless regions. It represents the Absolute Percentage Error APE  $(\hat{y}, y) = |\hat{y} - y/y|$  between the predicted log-radiance gradient  $\frac{\partial}{\partial t} \log \hat{L}(\boldsymbol{u}, t)$ , computed using auto-differentiation, and the finite difference approximation of the target log-radiance gradient  $\frac{\partial}{\partial t} \log L(\boldsymbol{u}, t) \approx \frac{pC_p}{t_{curr} - t_{ref}}$ , at a timestamp  $t_{sam}$  sampled between  $t_{ref}$  and  $t_{curr}$ , as follows:

$$\ell_{grad}(\boldsymbol{e}) = \operatorname{APE}\left(\frac{\partial}{\partial t}\log \hat{L}(\boldsymbol{u}, t_{sam}), \frac{pC_p}{t_{curr} - t_{ref}}\right) .$$
(4)

**Gamma Correction.** Since event cameras mainly provide observations of *changes* in log-radiance, not *absolute* log-radiance, the predicted log-radiance log  $\hat{L}$  from the reconstructed NeRF is only accurate up to an offset per color channel. There will be an additional channel-consistent scale ambiguity, when only the *Contrast Threshold Ratio*  $C_{\pm 1}/C_{-1}$  is known during reconstruction. Nonetheless, these ambiguities can be resolved post-reconstruction, given a set of reference images. Specifically, *ordinary least squares* can be used to perform an *affine* correction on log  $\hat{L}$ , or equivalently a *gamma* correction on  $\hat{L}$ , as follows:

$$\log \hat{\boldsymbol{L}}_{corr} = a \odot \log \hat{\boldsymbol{L}} + \boldsymbol{b} , \qquad (5)$$

where a and b are the correction parameters.



**Fig. 2:** Robust *e*-NeRF event generation model [27]

**Fig. 3:** Core analog circuit of a typical event sensor pixel (Adapted from [32])

# 4 Our Method

We first introduce the physically-accurate pixel bandwidth model proposed to model the motion blur of events (Sec. 4.1), which extends the event generation model of Robust *e*-NeRF (Sec. 3). Subsequently, we detail how to synthesize motion-blurred (effective) log-radiance incident at a pixel (Sec. 4.2), thus event motion blur, for optimization of a blur-minimal NeRF from motion-blurred events (Sec. 4.3). Lastly, we present an enhanced variant of gamma correction (Sec. 3) that takes the pixel bandwidth model into consideration (Sec. 4.4).

#### 4.1 Pixel Bandwidth Model

Event motion blur is attributed to the limited bandwidth of the sensor pixel analog circuit, which also bounds the minimum event detection latency, maximum frequency of change detectable and hence maximum event generation rate.

**Event Pixel Circuit.** Fig. 3 shows the core analog circuit of a typical event sensor pixel (in particular, that of the DVS128 [24]), which consists of 4 stages.

The Logarithmic Photoreceptor (Stage A) contains a Photodiode (PD) that transduces radiance (more accurately, irradiance) incident at the pixel to Signal Photocurrent  $I_p$  proportionally, and an active feedback loop that outputs a voltage  $V_p$  proportional to the logarithm of the Photocurrent  $I = I_p + I_{dark}$  [10,24,32]. A small Dark Current  $I_{dark}$  flows through the photodiode, even in the dark.

Next,  $V_p$  is buffered with a *Source Follower* (Stage B) to isolate the sensitive photoreceptor from rapid transients in subsequent stages [24,32]. The *Differencing Amplifier* (Stage C) then amplifies the change in source follower buffer output  $V_{sf}$  from the reset/reference voltage level, & outputs a voltage  $V_{diff}$  to be compared with both ON & OFF thresholds for event detection (Stage D) [10,24,32]. When  $V_{diff}$  exceeds either thresholds, an event is generated and the differencing amplifier is held in reset for a duration of the refractory period  $\tau$  [10,24,32].

**Model.** The design of the pixel analog circuit entails that event cameras in fact respond to changes in *effective log-radiance*  $\log L = \log (L_{sig} + L_{dark})$  instead,

where  $L_{sig}$  is the actual incident radiance signal and  $L_{dark}$  is the black level. Similar to standard image sensors, the black level is defined as the dark currentequivalent incident radiance, which is exponentially sensitive to temperature [32] and effectively limits the dynamic range of the sensor [10].

More precisely, as the pixel bandwidth is mainly limited by the first 3 stages of the analog circuit [3,5,8,9,23,30], event cameras effectively measure changes in the *low-pass-filtered/motion-blurred effective log-radiance* log  $L_{blur}$ . This explains the discussed motion blur of events and event sensor dynamic performance limit.

We accurately model the band-limiting behavior of the pixel with a unitygain 4<sup>th</sup>-order Non-Linear Time-Invariant (NLTI) Low-Pass Filter (LPF) in state-space form, with input  $u = \log L$ , state  $\boldsymbol{x} = \begin{bmatrix} \partial \log L_p / \partial t \log L_p \log L_{sf} \log L_{diff} \end{bmatrix}^{\top}$ and output  $\boldsymbol{y} = \begin{bmatrix} \log L_{sf} \log L_{diff} \end{bmatrix}^{\top}$ , as follows:

$$\dot{\boldsymbol{x}}(t) = A\left(\boldsymbol{u}(t)\right) \ \boldsymbol{x}(t) + B\left(\boldsymbol{u}(t)\right) \ \boldsymbol{u}(t) \\ \boldsymbol{y}(t) = C \ \boldsymbol{x}(t)$$
(6)

where 
$$A(u) = \begin{bmatrix} -2\zeta(u)\omega_n(u) & -\omega_n^2(u) & 0 & 0\\ 1 & 0 & 0 & 0\\ 0 & \omega_{c,sf} & -\omega_{c,sf} & 0\\ 0 & 0 & \omega_{c,diff} & -\omega_{c,diff} \end{bmatrix}, \ B(u) = \begin{bmatrix} \omega_n^2(u) \\ 0 \\ 0 \\ 0 \end{bmatrix},$$
$$C = \begin{bmatrix} 0 & 0 & 1 & 0\\ 0 & 0 & 1 \end{bmatrix}.$$

Specifically, the pixel bandwidth model is formed by a cascade of:

- 1. A unity-gain 2<sup>nd</sup>-order NLTI LPF, with input log L, state  $\left[ \frac{\partial \log L_p}{\partial t \log L_p} \right]^{\perp}$ and output log  $L_p$ , that models the transient response of the logarithmic photoreceptor [3,5,8,9,23,24,30]. Similar to its *Linear Time Invariant* (LTI) counterpart, this 2<sup>nd</sup>-order filter is characterized by its *Damping Ratio*  $\zeta$  and *Natural Angular Frequency*  $\omega_n$ . However, they are not constants, but complex non-linear functions of its input [3,5,23,24] (more details in the supplement). The bandwidth of this filter is mostly proportional to the exponential of its input exp u = L, which explains the susceptibility to motion blur under low-light. However, black level  $L_{dark}$  limits the minimum pixel bandwidth.
- 2. A unity-gain 1<sup>st</sup>-order LTI LPF, with input log  $L_p$  and state/output log  $L_{sf}$ , that models the transient response of the source follower buffer [9,23,30]. It is characterized by its constant bandwidth/*Cutoff Angular Frequency*  $\omega_{c,sf}$ , that is proportional to the source follower buffer bias current  $I_{sf}$  (Fig. 3) [23].
- 3. Another unity-gain 1<sup>st</sup>-order LTI LPF, with input log  $L_{sf}$ , state/output log  $L_{diff}$  and cutoff angular frequency  $\omega_{c,diff} > \omega_{c,sf}$  [23], that models the transient response of the differencing amplifier [23,30].

We also model the steady-state behavior of the differencing amplifier reset mechanism as a reset of the amplifier LPF state/output log  $L_{diff}$  to its input log  $L_{sf}$ , at the end of the refractory period (*i.e.* reference timestamp  $t_{ref}$ ). These 2 models allow the motion-blurred effective log-radiance log  $L_{blur}$  to be derived as:

 $\log L_{blur}(t) = \log L_{diff}(t) + \log L_{delta}(t_{ref}) \ e^{-\omega_{c,diff}(t-t_{ref})} \ , \ t \ge t_{ref} \ , \tag{7}$ 

where  $\log L_{delta} = \log L_{sf} - \log L_{diff}$ .

### 4.2 Synthesis of Motion-Blurred Effective Log-Radiance

The pixel bandwidth model proposed in Sec. 4.1 provides a means to accurately synthesize motion-blurred effective log-radiance  $\log L_{blur}$ , thus simulate event motion blur, given the pixel-incident log-radiance signal  $\log L_{sig}$ . However, the continuous-time and non-linear nature of the model (Eq. (6)) prohibits its direct computational implementation, for use in a simulator or for NeRF reconstruction using an Analysis-by-Synthesis framework (Sec. 4.3). A discrete-time counterpart of the model that operates on discrete-time input samples is necessary.

**Discrete-Time Model.** Assume the discrete-time sequence of inputs  $u[k] = \log L[k]$  is sampled at timestamps  $t_k$ , where the time intervals between successive samples  $\delta t_k = t_{k+1} - t_k$  may possibly be irregular. We derive a discrete-time model from the continuous-time, non-linear pixel bandwidth model by first linearizing [6] the 4<sup>th</sup>-order NLTI LPF (Eq. (6)) at different steady-state operating points  $(\bar{\boldsymbol{x}}[k], \bar{\boldsymbol{u}}[k]) = ([0 \ u[k+1] \ u[k+1] \ u[k+1]]^{\top}, u[k+1]),$  for each time interval  $(t_k, t_{k+1}]$ . Then, we discretize the linearized model assuming *First-Order Hold* (FOH) [6] (*i.e.* piecewise-linear) inputs *u*. This yields a discrete-time 4<sup>th</sup>-order *Linear Time-Varying* (LTV) LPF, in non-standard state-space form, as follows:

$$\begin{aligned} \boldsymbol{x}[k+1] &= A_d[k] \; \boldsymbol{x}[k] + B_d[k] \; u[k] + B_d[k] \; u[k+1] \\ \boldsymbol{y}[k] &= C_d \; \boldsymbol{x}[k] \end{aligned}$$
(8)

where  $A_d[k] = \Phi[k], \ B_d[k] = \Gamma_1[k] - \Gamma_2[k], \ \widetilde{B}_d[k] = \Gamma_2[k], \ C_d = C$  and:

$$\begin{bmatrix} \Phi[k] \ \Gamma_1[k] \ \Gamma_2[k] \\ 0 \ I \ I \\ 0 \ 0 \ I \end{bmatrix} = \exp\left( \begin{bmatrix} A \left( u[k+1] \right) \delta t_k \ B \left( u[k+1] \right) \delta t_k \ 0 \\ 0 \ 0 \ I \end{bmatrix} \right)$$

Numerical Solution. The discrete-time model presented can be directly integrated in an existing event simulator, e.g. ESIM [37] & its improved variant introduced in Robust e-NeRF, to synthesize  $\log L_{blur}$  & thus simulate event motion blur (more details in the supplement), assuming some appropriate initial state  $\boldsymbol{x}[k_0]$ , e.g. the steady-state on the initial input  $\boldsymbol{u}[k_0]$ . However, this cannot be done when the appropriate  $\boldsymbol{x}[k_0]$  is not well defined for an arbitrary  $\log L_{blur}[k]$ of interest, e.g. during NeRF reconstruction. We tackle this issue with the (numerical) solution to the transient response of the discrete-time model below:

$$\boldsymbol{y}[k] = C_d \left[ \varphi(k_0, k) \; \boldsymbol{x}[k_0] + \sum_{i=k_0}^{k-1} \varphi(i+1, k) \left( B_d[i] \; u[i] + \widetilde{B}_d[i] \; u[i+1] \right) \right] \;, \; (9)$$

#### 8 W. F. Low and G. H. Lee

where the state transition matrix  $\varphi(m,n) = \prod_{j=1}^{n-m} A_d[n-j]$ .

As the linearized, thus discretized, model is asymptotically stable, the magnitude of eigenvalues of  $A_d[k]$ , for all k, must be smaller than 1. This entails that  $\lim_{k \to \infty} \varphi(k_0, k) = 0$ . Thus, for a sufficiently long numerical integration time interval  $(t_{k_0}, t_k]$ ,  $\boldsymbol{y}[k]$  can be approximated as the zero-state response of the model, which is just a weighted sum of past & present inputs u between  $k_0 \& k$ :

$$\boldsymbol{y}[k] \approx \sum_{i=k_0}^{k} \boldsymbol{w}[i] \ u[i] \approx \sum_{i=k_0}^{k} \hat{\boldsymbol{w}}[i] \ u[i] \ , \tag{10}$$

where 
$$\boldsymbol{w}[i] = \begin{cases} C_d \varphi(k_0+1,k) B_d[k_0] & \text{, if } i = k_0 \\ C_d \left( \varphi(i+1,k) B_d[i] + \varphi(i,k) \widetilde{B}_d[i-1] \right) & \text{, if } i = k_0 + 1, \dots, k-1 \\ C_d \widetilde{B}_d[k-1] & \text{, if } i = k \end{cases}$$
  
 $\hat{\boldsymbol{w}}[i] = \boldsymbol{w}[i] \oslash \sum_{j=k_0}^k \boldsymbol{w}[j] ,$ 

and  $\oslash$  denotes *Hadamard*/element-wise division. It can be shown that  $\lim_{k \to k_0 \to \infty} \sum_{i=k_0}^{k} \boldsymbol{w}[i] = 1$ . Thus, we use the sum-normalized weights  $\hat{\boldsymbol{w}}[i]$  in practice, as they are corrected for the bias due to a finite integration interval  $(t_{k_0}, t_k]$ . In general, these weights tend to be larger for larger input (*i.e.* higher effective log-radiance) samples with timestamps closer to the output timestamp  $t_k$ .

**Importance Sampling.** Often times, we are interested in computing  $\log L_{blur}$  at some desired timestamp  $t_k$ , given only a function for sampling inputs  $u = \log L$  and a fixed sampling budget. To this end, we propose to infer the optimal input sample timestamps, represented by the random variable  $T_i \in (-\infty, t_k]$ , by sampling them from the following transformed exponential distribution:

$$T_i \sim \operatorname{Exp}\left(t_k - t_i; \ \omega_{c,dom,min}\right) = \omega_{c,dom,min} \ e^{-\omega_{c,dom,min}(t_k - t_i)} \ , \qquad (11)$$

where  $\omega_{c,dom,min}$  is the minimum possible dominant cutoff angular frequency of the pixel, achieved under extreme low-light when  $L = L_{dark}$ .

The suggested proposal distribution coarsely approximates the distribution represented by  $\hat{w}[i]$  over the interval  $(t_{k_0}, t_k]$ , as it corresponds to the weight function, derived from the zero-state response of the continuous-time dominant pole-approximated model under extreme low-light. Thus, it generally concentrate samples at relevant parts of the input (*i.e.* with large weight), achieving a similar goal as *importance sampling*. It works best under low-light, when event cameras are most susceptible to motion blur. More details are available in the supplement.

#### 4.3 Training

We employ the same training procedure used in Robust e-NeRF (Sec. 3) to optimize a blur-minimal NeRF from motion-blurred events in an Analysis-by-Synthesis manner, with a few exceptions. Specifically, we apply our training losses on the predicted motion-blurred effective log-radiance  $\log \hat{L}_{blur}$ , which was synthesized from NeRF renders  $\hat{L}$  using the proposed numerical solution to the pixel bandwidth model (Eqs. (7) and (10)) with importance sampling (Eq. (11)).

Moreover, we adopt a Huber-norm ( $\delta = 1$ ) variant of the threshold-normalized difference loss  $\ell_{diff}$  (Eq. (3)), which is less sensitive to outliers, with weight  $\lambda_{diff}$ . We also propose the threshold-normalized total variation loss  $\ell_{tv}$ , as a replacement for the target-normalized gradient loss  $\ell_{qrad}$  (Eq. (4)), with weight  $\lambda_{tv}$ .

**Fundamental Limitations.** When the black level  $L_{dark}$  is unknown, the incident radiance signal  $L_{sig}$  and  $L_{dark}$  cannot be unambiguously disentangled from just the observation of effective radiance  $L = L_{sig} + L_{dark}$  given by events. Thus, we can only reconstruct a NeRF with volume renders  $\hat{L}$  that represent predicted effective radiance, not predicted incident radiance signal as suggested by Robust *e*-NeRF. This is enabled by the assumption that the temperature of the event camera remains effectively constant over the entire duration of the given event steam, so that the dark current & thus black level remains effectively stationary.

Furthermore, as the pixel bandwidth depends on the *absolute* effective radiance L, the predicted effective radiance  $\hat{L}$  is theoretically gamma-accurate (*i.e.* gamma correction of  $\hat{L}$  is unnecessary), assuming known pixel bandwidth model parameters, contrary to what is suggested by Robust *e*-NeRF. However, in practice,  $\hat{L}$  is generally only gamma-accurate if the L associated to the events has a significant impact on the pixel transient response. In other words, the gamma-accuracy of  $\hat{L}$  greatly depends on the severity of event motion blur, hence camera speed, scene illumination, scene texture complexity, camera used and its settings.

**Threshold-Normalized Total Variation Loss.** This loss penalizes the mean contrast threshold  $\bar{C} = \frac{1}{2}(C_{-1}+C_{+1})$  normalized total variation of the predicted motion-blurred effective log-radiance log  $\hat{L}_{blur}$ , on a subinterval  $(t_{start}, t_{end}]$  sampled between the interval  $(t_{ref}, t_{curr}]$  given by an event, as follows:

$$\ell_{tv}(\boldsymbol{e}) = \left| \frac{\delta \log \hat{L}_{blur}}{\bar{C}} \right| \,, \tag{12}$$

where  $\delta \log \hat{L}_{blur} := \log \hat{L}_{blur}(\boldsymbol{u}, t_{end}) - \log \hat{L}_{blur}(\boldsymbol{u}, t_{start}).$ 

Similar to  $\ell_{grad}$  (Eq. (4)), this loss acts as a smoothness constraint for regularization of textureless regions in the scene. However, it imposes a stronger bias to enforce the uniformity of  $\log \hat{L}_{blur}$  between event intervals, which greatly helps with reconstructing large uniform patches. It can also effectively generalize to arbitrary threshold values due to the normalization, similar to  $\ell_{diff}$  (Eq. (3)).

Joint Optimization of Pixel Bandwidth Model Parameters. Our method does not strictly rely on the *pixel bandwidth model parameters*  $\Omega$  to be known as *a priori*, as it generally supports their joint optimization with *NeRF parameters*  $\Theta$ . However, their prior knowledge generally facilitates a more accurate reconstruction. The parameters  $\Omega$ , which include the parameters of non-linear damping ratio function  $\zeta$  & natural angular frequency function  $\omega_n$ , and angular cutoff frequencies  $\omega_{c,sf}$  and  $\omega_{c,diff}$  (Eq. (6)), depend on the pixel circuit design, semiconductor manufacturing process & user-defined event camera bias settings.

### 4.4 Translated-Gamma Correction

To eliminate the unknown black level offset and resolve potential gamma-inaccuracies in the predicted effective radiance  $\hat{L}$ , we propose a *Translated-Gamma Correction* on  $\hat{L}$  post-reconstruction, using a set of reference images, as follows:

$$\hat{\boldsymbol{L}}_{siq,corr} = \boldsymbol{b} \odot \hat{\boldsymbol{L}}^a - \boldsymbol{c} , \qquad (13)$$

where a, b and c are the correction parameters, via *Levenberg-Marquardt* nonlinear least squares optimization. The translation/offset correction is done independently per color channel to account for channel-varying spectral sensitivities.

### 5 Experiments

We conduct a series of *novel view synthesis* experiments, both on synthetic (Sec. 5.1) and real event sequences (Sec. 5.2), to verify that our method, Deblur *e*-NeRF, can indeed directly and effectively reconstruct blur-minimal NeRFs from motion-blurred events, generated under high-speed or low-light conditions, using a physically-accurate pixel bandwidth model.

Metrics. We adopt the commonly used PSNR, SSIM [45] and AlexNet-based LPIPS [52] to evaluate the similarity between the target and translated-gamma-corrected (Sec. 4.4) synthesized novel views, for all methods in each experiment.

**Baselines.** We benchmark our method, Deblur *e*-NeRF, against the state-ofthe-art, Robust *e*-NeRF [27], and a naïve baseline of E2VID [38] (a seminal events-to-video reconstruction method) + NeRF [31] (as well as 2 other *image* blur-aware baselines in the supplement). The setup of an *event* motion bluraware baseline is hindered by the lack of relevant works. The implementation of all methods employ a common NeRF backbone [22] to allow for a fair comparison.

**Datasets.** We perform the synthetic experiments on the default set of sequences released by Robust *e*-NeRF, and a novel set similarly simulated on the "Realistic Synthetic 360°" scenes [31]. However, our sequences involve event motion blur due to fast camera motion and/or poor scene illumination. Such sequences were simulated with our proposed event simulator, which incorporates the pixel bandwidth model presented in Secs. 4.1 and 4.2 in the improved ESIM [37] event simulator introduced in Robust *e*-NeRF. Similar to the sequences in the Robust e-NeRF synthetic event dataset, the events are generated from a virtual event camera moving in a hemi-/spherical spiral motion about the object at the origin.

On the other hand, we conduct the real experiments on the 08\_peanuts\_running and 11\_all\_characters sequences of the EDS dataset [11], which are generally 360° captures of a table of objects in an office room under moderate lighting. These sequences were chosen as they involve some high-speed camera motion.



Fig. 4: Pixel bandwidth of DVS128 [24] with nominal biases

#### 5.1 Synthetic Experiments

The synthetic experiments serve as the main benchmark to assess all methods, as they enable controlled tests under diverse realistic conditions with precise ground truth, a task that would otherwise be infeasible using real sequences.

As with the default sequences in Robust e-NeRF [27], we simulate ours with symmetric contrast thresholds of 0.25 (*i.e.*  $C_{-1} = C_{+1} = 0.25$ ), zero pixelto-pixel threshold standard deviation and refractory period (*i.e.*  $\sigma_{C_p} = 0, \tau =$ 0), and provide camera poses at 1 kHz. Moreover, the pixel bandwidth model parameters that we adopt correspond to that of the DVS128 [24] event camera with nominal biases (provided in jAER [4]). Fig. 4 illustrates the light-dependent behaviour of its pixel bandwidth, which is mostly proportional to the scene illuminance  $E_{sc}$ , thus radiance incident at a pixel  $L_{sig}$ . Unless otherwise stated, the virtual event camera revolves the object 4 times with uniform 2 revolution per second speed about the object vertical axis. All sequences are also simulated under a scene illuminance  $E_{sc}$  of 1 000lux by default, which corresponds to standard office lighting [3,5]. Under such a condition, the pixel bandwidth spans around 50–2500Hz, depending on the incident radiance. Due to limited resources, our method is only trained with  $\frac{1}{8\times}$  the batch size of our baselines, by default.

**Upper Bound Performance.** To quantify the upper bound performance, we evaluate all methods on motion blur-free event sequences, which is effectively given by a pixel with infinite bandwidth. For this purpose, we remove the pixel bandwidth model from our method and train it with the same batch size as the baselines. The main difference between Robust *e*-NeRF and our method, under such a setting, is the replacement of the target-normalized gradient loss  $\ell_{grad}$  with the threshold-normalized total variation loss  $\ell_{tv}$ . Thus, the quantitative results reported in Tab. 1 verifies the effectiveness of our proposed  $\ell_{tv}$  over  $\ell_{grad}$ . Moreover, the qualitative results shown in Fig. 5, particularly at the back of the chair, clearly shows the strength of  $\ell_{tv}$  in regularizing large textureless patches.

**Table 1:** Upper bound performance without event motion blur

 Table 2: Quantitative results of the real exps.

					08_p	eanuts_ru	nning	11_all_characters		
Method	$\mathrm{PSNR}\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Method	$\mathrm{PSNR} \uparrow$	SSIM $\uparrow$	$_{\rm LPIPS}\downarrow$	$PSNR \uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
E2VID + NeRF	19.49	0.847	0.268	E2VID + NeRF	14.85	0.690	0.595	13.12	0.695	0.627
Robust e-NeRF	28.48	0.944	0.054	Robust e-NeRF	18.00	0.677	0.507	15.91	0.677	0.552
Deblur e-NeRF	29.43	0.953	0.043	Deblur $e$ -NeRF	18.27	0.695	0.503	16.53	0.710	0.511

**Table 3:** Effect of camera speed. <sup>†</sup>Trained with  $1/8 \times$  the batch size of baselines.

Method		v = 0.125	×		$v = 1 \times$		$v = 4 \times$			
	PSNR ·	$\uparrow$ SSIM $\uparrow$	LPIPS $\downarrow$	$PSNR \uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	$PSNR \uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	
E2VID + NeRF	18.58	0.849	0.259	18.85	0.839	0.278	17.82	0.804	0.328	
Robust $e$ -NeRF	28.31	0.943	0.050	26.11	0.924	0.074	22.18	0.861	0.122	
Deblur $e$ -NeRF <sup>†</sup>	28.71	0.948	0.048	28.41	0.947	0.049	27.48	0.939	0.061	

**Table 4:** Effect of scene illuminance. <sup>†</sup>Trained with  $1/8 \times$  the batch size of baselines.

Method	$E_{sc}$	$= 100 \ 00$	)0lux	$E_{so}$	= 1  000	)lux	$E_{sc} = 10 lux$			
	$\rm PSNR\uparrow$	SSIM $\uparrow$	$\text{LPIPS} \downarrow$	$PSNR\uparrow$	SSIM $\uparrow$	$\text{LPIPS} \downarrow$	$\mathrm{PSNR}\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	
E2VID + NeRF	19.27	0.846	0.268	18.85	0.839	0.278	17.24	0.804	0.354	
Robust $e$ -NeRF	27.62	0.942	0.055	26.11	0.924	0.074	22.72	0.870	0.129	
Deblur $e$ -NeRF <sup>†</sup>	28.73	0.948	0.047	28.41	0.947	0.049	28.62	0.935	0.059	

**Table 5:** Collective effect of camera speed and scene illuminance. <sup>†</sup>Trained with  $1/8 \times$  the batch size of baselines.

	Opt	$\mathcal{O}_{\mathrm{pt.}}$	$v=0.125\times,E_{sc}=100\;000 lux$			$v=1\times,E_{sc}=1\;000 \textit{lux}$			$v = 4 \times, E_{sc} = 10 lux$		
Method	$C_p \& \tau$		$PSNR \uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	$PSNR\uparrow$	SSIM $\uparrow$	$_{\rm LPIPS}\downarrow$	$PSNR \uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
E2VID + NeRF	_	_	19.19	0.844	0.281	18.85	0.839	0.278	15.37	0.799	0.436
Robust e-NeRF	×	_	28.27	0.944	0.057	26.11	0.924	0.074	18.42	0.814	0.255
	$\checkmark$	_	28.28	0.944	0.051	26.31	0.923	0.075	18.51	0.812	0.254
Deblur $e$ -NeRF <sup>†</sup>	×	×	29.00	0.950	0.043	28.41	0.947	0.049	26.15	0.904	0.134
	×	$\checkmark$	28.19	0.943	0.046	26.07	0.930	0.067	25.59	0.896	0.156

Effect of Camera Speed. To investigate the effect of event motion blur due to high-speed camera motion, we evaluate all methods on 3 sets of sequences simulated with camera speeds v that are  $0.125 \times$ ,  $1 \times \& 4 \times$  of the default setting, respectively. As event motion blur may lead to a significant "loss or introduction of events" (Sec. 1), we also quantify the average number of events relative to that of its corresponding blur-free sequence, across all sequences in the set. This translates to 93.36%, 100.95% & 95.21% for  $v = 0.125 \times$ ,  $1 \times \& 4 \times$ , respectively.

The quantitative results in Tab. 3 clearly underscores the significance of incorporating a pixel bandwidth model, as our method significantly outperforms all baselines, especially under high-speed motion, despite being trained with  $1/8\times$  the batch size of baselines. The results also display our robustness to event motion blur, as our performance remains relatively unperturbed under varying camera speeds, & remains close to our upper bound performance. Qualitative results in Fig. 5 further validate our effectiveness in reconstructing a blur-minimal NeRF, as our method is free from artifacts such as floaters and double edges.

Effect of Scene Illuminance. To assess the effect of event motion blur due to poor scene illumination, we benchmark all methods on 3 sets of sequences simulated with scene illuminance  $E_{sc}$  of 100 000, 1 000 & 10*lux*, which correspond to sunlight, office light & street light, respectively [3]. Furthermore, the pixel bandwidth spans around 1200–4000*Hz*, 50–2500*Hz* & 20–170*Hz*, respectively, depending on the incident radiance. They also have 82.49%, 100.95% & 17.22% the number of events of the blur-free set, respectively. The quantitative results given in Tab. 4 undoubtedly verifies the importance of modeling event motion blur, as our method once again outperform other works, while remaining close to the upper bound performance. It also reveals our astonishing robustness to blur-induced "loss of events", particularly under  $E_{sc} = 10lux$ .

Collective Effect. To assess the effect of event motion blur due to both high speed and low light, which resemble challenging real-world conditions, we benchmark all methods on 3 sets of sequences with different difficulty levels: easy  $(v = 0.125 \times, E_{sc} = 100\ 000 lux)$ , medium  $(v = 1 \times, E_{sc} = 1\ 000 lux)$  and hard  $(v = 4 \times, E_{sc} = 10 lux)$ , which have 93.36%, 100.95% and 7.14% the number of events of the blur-free set, respectively. We also evaluate Robust *e*-NeRF with jointly optimized contrast thresholds  $C_p$  & refractory period  $\tau$  to validate the importance of modeling event motion blur. We benchmark our method with jointly optimized pixel bandwidth model parameters  $\Omega$ , poorly initialized to  $4 \times$  their true value, to assess the robustness of joint optimization, thus auto-calibration.

The quantitative and qualitative results given in Tab. 5 and Fig. 5, respectively, generally reflect the results of the 2 previous experiments, thus a similar conclusion can be drawn. In addition,  $C_p \& \tau$  clearly cannot compensate for the lack of a pixel bandwidth model, as their joint optimization has virtually no effect on Robust *e*-NeRF. The quantitative results also reveal the feasibility of jointly optimizing  $\Omega$ , especially under severe motion blur. While our performance deteriorates as the difficulty increases, which suggests a limit to our robustness to event motion blur, we show in the supplement that increasing our batch size to that of the baselines significantly improves our performance, hence robustness.

#### 5.2 Real Experiments

To account for unknown event camera intrinsic parameters, we train Robust *e*-NeRF and our method with jointly optimized contrast thresholds and refractory period. We also train our method with jointly optimized pixel bandwidth model parameters, initialized from DVS128 [24] fast biases (provided in jAER [4]), using the same batch size as the baselines. While the 2 sequences only occasionally involve fast camera motion under moderate office lighting, the quantitative results reported in Tab. 2 demonstrate our superior performance. This is also supported qualitatively in Fig. 5, where the table & walls are visibly more uniform, and less blooming artifacts are observed around objects, compared to Robust e-NeRF.



Fig. 5: Qualitative results of synthetic & real experiments, w/o jointly optimized contrast thresholds, refractory period and pixel bandwidth model parameters, if applicable

# 6 Conclusion

In this paper, we introduce Deblur *e*-NeRF, a novel method to directly and effectively reconstruct blur-minimal NeRFs from motion-blurred events, generated under high-speed or low-light conditions. The core component of this work is a physically-accurate pixel bandwidth model that accounts for event motion blur. We also propose a threshold-normalized total variation loss to better regularize large textureless patches. Despite its accomplishments, Deblur *e*-NeRF still inherits the limitations of Robust *e*-NeRF and other works, *e.g.* assumption of known camera trajectory. Moreover, since the synthesis of motion-blurred effective log-radiance at a given timestamp requires multiple past and present samples of effective log-radiance, given by NeRF renders, our reconstruction also incurs a higher computational and memory cost. Joint optimization of pixel bandwidth model parameters is also sometimes unstable. We leave these as future work.

Acknowledgements. This research / project is supported by the National Research Foundation, Singapore, under its NRF-Investigatorship Programme (Award ID. NRF-NRFI09-0008), and the Tier 1 grant T1-251RES2305 from the Singapore Ministry of Education.

# References

- Cannici, M., Scaramuzza, D.: Mitigating motion blur in neural radiance fields with events and frames. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- 2. Cho, S., Lee, S.: Fast motion deblurring. ACM Trans. Graph. (2009)
- Delbruck, T., Mead, C.A.: Analog VLSI Phototransduction by continuous-time, adaptive, logarithmic photoreceptor circuits. Vision Chips: Implementing vision algorithms with analog VLSI circuits (1995)
- 4. Delbruck, T.: Frame-free dynamic digital vision. In: Proceedings of Intl. Symp. on Secure-Life Electronics, Advanced Electronics for Quality Life and Society (2008)
- 5. Delbruck, T.: Investigations of analog VLSI visual transduction and motion processing. Doctoral Thesis, California Institute of Technology (1993)
- Franklin, G., Powell, J., Workman, M.: Digital Control of Dynamic Systems. Addison-Wesley series in electrical and computer engineering: Control engineering, Addison-Wesley (1998)
- Gallego, G., Delbrück, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., Leutenegger, S., Davison, A.J., Conradt, J., Daniilidis, K., Scaramuzza, D.: Eventbased vision: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence (2020)
- 8. Graca, R., Delbruck, T.: Unraveling the paradox of intensity-dependent dvs pixel noise. In: International Image Sensor Workshop (IISW) (2021)
- Graca, R., McReynolds, B., Delbruck, T.: Optimal biasing and physical limits of dvs event noise. In: International Image Sensor Workshop (IISW) (2023)
- Graça, R., McReynolds, B., Delbruck, T.: Shining light on the dvs pixel: A tutorial and discussion about biasing and optimization. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) (2023)
- Hidalgo-Carrió, J., Gallego, G., Scaramuzza, D.: Event-Aided Direct Sparse Odometry. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Hu, Y., Liu, S.C., Delbruck, T.: v2e: From Video Frames to Realistic DVS Events. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2021)
- Hwang, I., Kim, J., Kim, Y.M.: Ev-nerf: Event based neural radiance field. In: IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2023)
- Jiang, Z., Zhang, Y., Zou, D., Ren, J., Lv, J., Liu, Y.: Learning event-based motion deblurring. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
- Joubert, D., Marcireau, A., Ralph, N., Jolley, A., van Schaik, A., Cohen, G.: Event Camera Simulator Improvements via Characterized Parameters. Frontiers in Neuroscience (2021)
- Kato, H., Beker, D., Morariu, M., Ando, T., Matsuoka, T., Kehl, W., Gaidon, A.: Differentiable Rendering: A Survey (2020), arXiv:2006.12057 [cs]
- 17. Klenk, S., Koestler, L., Scaramuzza, D., Cremers, D.: E-nerf: Neural radiance fields from a moving event camera. IEEE Robotics and Automation Letters (2023)
- Krishnan, D., Fergus, R.: Fast image deconvolution using hyper-laplacian priors. In: Proceedings of the 22nd International Conference on Neural Information Processing Systems (2009)

- 16 W. F. Low and G. H. Lee
- Krishnan, D., Tay, T., Fergus, R.: Blind deconvolution using a normalized sparsity measure. In: CVPR 2011 (2011)
- Lee, D., Lee, M., Shin, C., Lee, S.: Dp-nerf: Deblurred neural radiance field with physical scene priors. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- Lee, D., Oh, J., Rim, J., Cho, S., Lee, K.: Exblurf: Efficient radiance fields for extreme motion blurred images. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- Li, R., Gao, H., Tancik, M., Kanazawa, A.: Nerfacc: Efficient sampling accelerates nerfs. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- 23. Lichtsteiner, P.: An AER temporal contrast vision sensor. Doctoral Thesis, ETH Zurich (2006)
- 24. Lichtsteiner, P., Posch, C., Delbruck, T.: A  $128 \times 128$  120 dB 15  $\mu$ s latency asynchronous temporal contrast vision sensor. IEEE Journal of Solid-State Circuits (2008)
- Lin, S., Zhang, J., Pan, J., Jiang, Z., Zou, D., Wang, Y., Chen, J., Ren, J.: Learning event-driven video deblurring and interpolation. In: Computer Vision – ECCV 2020 (2020)
- Liu, H., Peng, S., Zhu, L., Chang, Y., Zhou, H., Yan, L.: Seeing motion at nighttime with an event camera. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Low, W.F., Lee, G.H.: Robust e-nerf: Nerf from sparse & noisy events under nonuniform motion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- Ma, L., Li, X., Liao, J., Zhang, Q., Wang, X., Wang, J., Sander, P.V.: Deblur-NeRF: Neural Radiance Fields From Blurry Images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Ma, Q., Paudel, D.P., Chhatkuli, A., Van Gool, L.: Deformable neural radiance fields using rgb and event cameras. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- 30. McReynolds, B.J., Graca, R.P., Delbruck, T.: Experimental methods to predict dynamic vision sensor event camera performance. Optical Engineering (2022)
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In: European Conference on Computer Vision (ECCV) 2020 (2020)
- 32. Nozaki, Y., Delbruck, T.: Temperature and Parasitic Photocurrent Effects in Dynamic Vision Sensors. IEEE Transactions on Electron Devices (2017)
- 33. Park, H., Mu Lee, K.: Joint estimation of camera pose, depth, deblurring, and super-resolution from a blurred image sequence. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2017)
- Perrone, D., Favaro, P.: Total variation blind deconvolution: The devil is in the details. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014)
- 35. Potmesil, M., Chakravarty, I.: Modeling motion blur in computer-generated images. SIGGRAPH Comput. Graph. (1983)
- Qi, Y., Zhu, L., Zhang, Y., Li, J.: E2nerf: Event enhanced neural radiance fields from blurry images. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
- Rebecq, H., Gehrig, D., Scaramuzza, D.: ESIM: an Open Event Camera Simulator. In: Proceedings of The 2nd Conference on Robot Learning, PMLR (2018)

- Rebecq, H., Ranftl, R., Koltun, V., Scaramuzza, D.: High speed and high dynamic range video with an event camera. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021)
- Rudnev, V., Elgharib, M., Theobalt, C., Golyanik, V.: Eventnerf: Neural radiance fields from a single colour event camera (2022), arXiv:2206.11896 [cs]
- Shan, Q., Jia, J., Agarwala, A.: High-quality motion deblurring from a single image. ACM Trans. Graph. (2008)
- 41. Son, H., Lee, J., Lee, J., Cho, S., Lee, S.: Recurrent video deblurring with blurinvariant motion estimation and pixel volumes. ACM Trans. Graph. (2021)
- Tao, X., Gao, H., Shen, X., Wang, J., Jia, J.: Scale-recurrent network for deep image deblurring. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (2018)
- Wang, P., Zhao, L., Ma, R., Liu, P.: Bad-nerf: Bundle adjusted deblur neural radiance fields. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- 44. Wang, Z., Cun, X., Bao, J., Zhou, W., Liu, J., Li, H.: Uformer: A general u-shaped transformer for image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: from error visibility to structural similarity. IEEE Transactions on Image Processing (2004)
- 46. Xu, F., Yu, L., Wang, B., Yang, W., Xia, G.S., Jia, X., Qiao, Z., Liu, J.: Motion deblurring with real events. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2583–2592 (October 2021)
- 47. Xu, L., Jia, J.: Two-phase kernel estimation for robust motion deblurring. In: Proceedings of the 11th European Conference on Computer Vision: Part I (2010)
- Xu, L., Zheng, S., Jia, J.: Unnatural l0 sparse representation for natural image deblurring. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2013)
- 49. Yang, Y., Liang, J., Yu, B., Chen, Y., Ren, J.S., Shi, B.: Latency correction for event-guided deblurring and frame interpolation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
- Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Multi-stage progressive image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
- Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
- Zhang, X., Yu, L.: Unifying motion deblurring and frame interpolation with events. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)