

From Corners to Fiducial Tags: Revisiting Checkerboard Calibration for Event Cameras

Taehun Ryu, Changwoo Kang, and Kyungdon Joo*
UNIST

{taehunryu, kangchangwoo, kyungdon}@unist.ac.kr

Abstract

The conventional checkerboard-based calibration for standard cameras faces fundamental limitations when applied to bio-inspired event cameras. Specifically, this stems from two challenges: (i) Events are triggered asynchronously at different timestamps along motion trajectories. If we accumulate them directly on the image plane, it causes temporal misalignment and produces blurred edges. (ii) Checkerboard corners on event cameras show near-zero event occurrence at the corner itself. This hinders reliable corner localization and makes calibration difficult. To address these issues, we present a novel calibration framework that directly detects checkerboard corners from event data without learning-based grayscale image reconstruction. We first mathematically analyze the absence of events at corner points. Based on this fact, we then leverage edge-driven event cues to initialize corner positions. Using the near-zero event occurrence at checkerboard corners, we gradually refine the estimated corner toward low event-density regions, achieving sub-pixel accuracy. Furthermore, we extend the corner detection to fiducial markers such as AprilTag, resulting in reliable detection even under partial visibility or occlusion. Evaluations on self-collected and public data demonstrate reliable checkerboard corner detection and stable camera calibration. Additional information is available at the following link: <https://vision3d-lab.github.io/corner2tag/>.

1. Introduction

The event camera, a neuromorphic vision sensor, has attracted widespread attention for its robustness to motion blur and low-light conditions compared to conventional cameras [13, 24]. Consequently, event-based vision has shown significant effectiveness in diverse computer vision and robotics domains such as 3D scanning [33], robotic manufacturing [1], high-speed object detection and tracking [17, 31, 34, 50], and autonomous navigation [8, 44].

*Corresponding author.




Figure 1. **Illustration of proposed pattern detection.** We utilize checkerboard patterns and AprilTag [46] to directly detect corners from an event stream (left). Our method can extract the corners of the checkerboard and the distinct fiducial (right).

As the demand for such applications increases, calibration of event cameras becomes crucial for reliable vision systems as a process of estimating camera parameters. For frame-based cameras, the camera parameters are typically determined by minimizing the reprojection error between 2D fiducial points detected from images and their corresponding 3D points pre-defined on the calibration target. Various pre-defined patterns (e.g., checkerboard [7], circular grid [22]) have been developed to provide such easily observable fiducial points for accurate calibration. Among them, the checkerboard pattern is the most common and reliable calibration target [3, 4]. Specifically, the fiducial points of the checkerboard are the corners of black-and-white squares, which are defined by the intersections of two edges. These characteristics maintain geometrically well-defined under perspective projection or lens distortion.

However, the checkerboard pattern remains unexplored for event cameras, as it cannot be directly applied to a raw event stream. This is because event cameras do not preserve the image structure. Instead, they record relative brightness changes asynchronously at each pixel as binary values. For this reason, some methods utilize learning-based reconstruction to convert a raw event stream into grayscale images, and then apply frame-based calibration to the reconstructed grayscale images [18, 35]. However, artifacts of reconstructed images inherently constrain calibration precision. Some methods utilize blinking LED boards to directly

process a raw event stream [19, 40]. They generally yield accurate calibration results, but they demand high hardware requirements and rely on dedicated procedures. Other approaches employ circle-grid patterns for event-based calibration [6, 23, 41]. They exhibit motion-invariant characteristics and achieve robust detection performance, yet their design inherits the drawback of circular patterns, which are sensitive to lens distortion [43]. Compared to the above approaches, the checkerboard pattern provides geometrically consistent and precisely localized corners, ensuring high calibration performance. It also holds strong potential for extension to ID-based markers such as ARTag [9], and AprilTag [36, 46], as they share a similar corner-based detection principle.

In this work, we advocate the revisit of the standard checkerboard pattern using event cameras to preserve its proven accuracy and its extensibility to fiducial markers (see Fig. 1). However, this approach faces two key challenges. First, due to the asynchronous characteristic of each event, stacking them into a single image produces blurred edges [11]. Second, events are not triggered at checkerboard corners due to their locally symmetric intensity pattern, which limits reliable corner detection.

To resolve these, we propose the first event camera calibration framework that directly detects checkerboard corners from event representation without grayscale image reconstruction (see Fig. 2). We create an Image of Warped Events (IWE) to temporally align asynchronous events. This process produces a sharp, static representation, which is essential for precise corner detection. Then, we provide a mathematical analysis of Event Rate (ER) at checkerboard corners. This analysis confirms that events are theoretically absent at the exact corner intersections. Grounded in this analysis, our method first initializes corners using strong edge cues, instead of finding corners directly. We then refine the initial corners by moving them toward the point of minimum event density to achieve sub-pixel accuracy. These accurate corner estimates are then utilized to optimize the camera parameters. Furthermore, we demonstrate that our framework extends to fiducial markers, such as AprilTag [46]. This enables reliable tag identification even under partial visibility.

Our main contributions are summarized as follows:

- **Event-based checkerboard calibration:** We present the first calibration framework using detected checkerboard corners without grayscale image reconstruction.
- **Mathematical analysis of corner events:** We analyze the absence of events at corner intersections, which we then exploit for corner initialization and refinement.
- **Analysis-based corner localization:** Based on our analysis of corner events, we initialize the corners and refine them to achieve sub-pixel accuracy.
- **Extension to fiducial markers:** Our framework gener-

alizes to fiducial markers (*e.g.*, AprilTag [46]), ensuring robust detection under partial visibility.

2. Related Work

Camera calibration process can be divided into two steps: Pattern detection and camera parameter optimization [48]. Pattern detection aims to identify accurate fiducial points that serve as 2D-3D correspondences from predefined calibration targets. Camera parameter optimization then estimates the intrinsic and extrinsic parameters of cameras by solving an optimization problem based on these correspondences. In this section, we mainly deal with pattern detection for both frame-based and event cameras. Please refer to other frameworks [3, 4, 10] for details of the optimization process.

Pattern Detection for Frame-Based Cameras. As fiducial points directly provide 2D-3D correspondences, various calibration targets have been proposed [4, 7, 22, 46]. Among them, the checkerboard pattern [3, 4, 7, 29] serves as the standard for high-precision calibration due to its geometrically well-defined and densely distributed corners, which enable robust and sub-pixel localization even under perspective transformation or lens distortion. Building upon this concept, ARTag [9], AprilTag [36, 46], and ArUco marker [14] encode unique marker IDs to ensure reliable detection under partial observations. Accordingly, the checkerboard remains the most widely adopted target in frameworks such as OpenCV [4]. In parallel, other primitives such as circular features or dots are used for calibration targets [22, 26, 27]. They use centroids of circular shapes as fiducial points. They exhibit rotational invariance but suffer from degraded localization accuracy under nonlinear transformations such as lens distortion [28, 43].

In this work, our objective is to apply the well-validated checkerboard pattern to event cameras while leveraging its proven calibration accuracy and structural extensibility to ID-coded fiducial systems [9, 14, 36, 46].

Pattern Detection for Event Cameras. Conventional pattern detection approaches are not directly applicable to event-based data due to its frameless and binary-valued nature [45]. As a straightforward approach, active-target (*e.g.*, LCD or blinking LED) and wand are initially adopted for event cameras [5, 19, 20, 40, 47]. Their design is well-tailored to the triggering mechanism of events, making them intuitive and accurate. However, these approaches require specialized hardware and corresponding methodologies.

In the other directions, several studies attempt to detect circular or circle-grid patterns directly from a raw event stream [6, 23, 41]. These patterns offer robustness, as events are generated along every edge regardless of motion direction. In these approaches, events are typically clustered and

modeled by fitting cylindrical or elliptical shapes to a raw event stream. They share the drawback of high sensitivity to lens distortion with circular patterns, with the cylindrical fitting making precise localization even more difficult.

With the recent advances in learning-based approaches, some methods [18, 35] leverage the event-to-video reconstruction network [39] to convert a raw event stream into grayscale images and apply the Harris corner detector [21] to identify fiducial points. However, event-to-video reconstruction introduces artifacts [49], which may degrade the accuracy of subsequent corner localization. These limitations motivate direct detection of fiducial points from event-domain representations rather than reconstructed images.

In this work, we detect fiducial points directly from event representation, specifically IWE [12], rather than relying on imprecise reconstruction. To this end, we analyze event generation at corner points and propose a calibration pipeline specialized to event characteristics.

3. Method

3.1. Overview

We investigate the event camera calibration directly in the event domain, without relying on grayscale image reconstruction. Calibrating event cameras with a checkerboard is challenging due to asynchronous event generation, which leads to blurred edges when events are accumulated [11]. In addition, the absence of events at the corner intersection makes corner detection difficult.

To address these, we propose an event camera calibration pipeline using a checkerboard pattern (see Fig. 2). We take as input a raw event stream $\mathcal{E} = \{e_k\}_{k=1}^N$, where each event $e_k = (\mathbf{x}_k, t_k, p_k)$ consists of the pixel location $\mathbf{x}_k$, timestamp t_k , and polarity p_k (-1 or $+1$). Then, we convert a raw event stream into an IWE [12] via events warping and events accumulation process to obtain sharper edges (Sec. 3.3). Given IWE, we initialize checkerboard corners by leveraging reliable edge cues, and then refine them toward the local minimum of event density to obtain sub-pixel accuracy. Both steps are motivated by our analysis in Sec. 3.2, which shows that events are stably generated along edges while vanishing at the corner. The refined corners are fed into the batch optimization process to estimate the camera parameters $\mathbf{K}, \mathbf{d}$ (Sec. 3.4). Moreover, we show that our framework generalizes to other fiducial markers, such as AprilTag [46], allowing for robust tag identification even under partial occlusion (Sec. 3.5).

3.2. Event Rate on Checkerboard Corner

Understanding how events are generated around an ideal checkerboard corner helps to explain the underlying characteristics of event data and to design more robust detection methods. Here, we analyze that events are not triggered in




Figure 2. **Overview of the proposed calibration pipeline.** The framework converts a raw event stream into both IWE and pIWE, which are representations well-suited for extracting corner points. These representations are then used by the pattern detection module, which feeds the detected corners into the intrinsic optimization process to estimate the final camera parameters $\mathbf{K}, \mathbf{d}$.

the checkerboard corner, using the intensity model of the checkerboard corner in frame-based cameras [29] and the event generation rate formulation [32]. Details of the analysis are described in the supplementary material.

Intensity Model on Checkerboard Corner. In frame-based cameras, the ideal intensity distribution of the checkerboard corner can be defined with a step function $\rho(\cdot)$ as [29]:

$$\hat{I}(\mathbf{x}) = \rho(\mathbf{n}_1^\top \mathbf{x})\rho(\mathbf{n}_2^\top \mathbf{x}), \quad \rho(x) = \begin{cases} +1 & \text{if } x > 0, \\ -1 & \text{otherwise,} \end{cases} \quad (1)$$

where $\mathbf{n}_i$ denotes unit normals to checkerboard edges. Using this definition, Krüger *et al.* [29] define the intensity model of the checkerboard corner $\tilde{I}$ by replacing the step function $\rho(\cdot)$ with a centered sigmoid smoothed by a Gaussian filter of width σ :

$$\tilde{I}(\mathbf{x}) \approx S(\mathbf{n}_1^\top \mathbf{x}, \sigma)S(\mathbf{n}_2^\top \mathbf{x}, \sigma), \quad (2)$$

where $S(x, \sigma) = 2/(1+e^{-x/\sigma})-1$ is a centered sigmoid. In this model, the ideal checkerboard corner is located $\mathbf{x} = \mathbf{0}$. Note that since construction $S(x, \sigma) \in (-1, 1)$, $\tilde{I}(\mathbf{x}) \in [-1, 1]$ represents the idealized contrast pattern of the checkerboard, rather than the actual brightness observed by the camera.

Event Generation Rate. Based on the event-triggering mechanism, an event is generated whenever the change in logarithmic intensity at a pixel exceeds a fixed contrast threshold C [13]. This implies that, at a given location, faster or larger brightness variations result in a greater number of events accumulated within a fixed time window. Thus, Event Rate (ER) can be defined as the number of events generated per unit time in a given time interval.

Therefore, ER can be interpreted as a statistical measure of the logarithmic intensity change rate, which is related to the image gradient [16, 32].

Following this interpretation, Lin *et al.* [32] establish a relationship between ER and the gradient value in short time interval $[t - \Delta t/2, t + \Delta t/2]$ as follows:

$$R_e(\mathbf{x}, t, \Delta t) \approx \frac{1}{C} |\nabla L(\mathbf{x}, t) \cdot \mathbf{v}|, \quad (3)$$

where $L(\mathbf{x}, t) = \log I(\mathbf{x}, t)$ denotes the logarithmic image intensity at pixel location $\mathbf{x}$, and $\mathbf{v}$ is the motion vector. For simplicity, we omit the temporal variable t in the following analysis, using $L(\mathbf{x})$ and $R_e(\mathbf{x})$ instead of $L(\mathbf{x}, t)$ and $R_e(\mathbf{x}, t, \Delta t)$, respectively.

Event Rate on Checkerboard Corner. To analyze ER on checkerboard corner, we combine Eq. (2) and Eq. (3). Before combining them, we first reformulate centered sigmoid $S(x, \sigma)$, noting that $S(x, \sigma)$ can be expressed as a hyperbolic tangent function,

$$S(x, \sigma) = \tanh\left(\frac{x}{2\sigma}\right). \quad (4)$$

Then, we recall that the hyperbolic tangent function has the following Taylor expansion [25]:

$$\tanh(x) = x - \frac{1}{3}x^3 + \frac{2}{15}x^5 - \dots, \text{ when } |x| < \frac{\pi}{2}. \quad (5)$$

Since we are primarily interested in the neighborhood of the ideal checkerboard corner, where its region is in $|x| \ll 1$, the Taylor expansion of $\tanh(x)$ is valid, and the higher-order terms can be neglected as

$$S(x, \sigma) = \tanh\left(\frac{x}{2\sigma}\right) \approx \frac{x}{2\sigma} + \mathcal{O}(x^3) \approx \frac{x}{2\sigma}, \quad (6)$$

where $\mathcal{O}(x^3)$ denotes the third-order remainder in the Taylor expansion.

Based on this convention, we reformulate Eq. (2) using the approximation in Eq. (6), deriving its local analytical form around the checkerboard corner:

$$\tilde{I}(\mathbf{x}) \approx \frac{1}{4\sigma^2} (\mathbf{n}_1^\top \mathbf{x})(\mathbf{n}_2^\top \mathbf{x}), \quad \nabla \tilde{I}(\mathbf{x}) \approx G \mathbf{x}, \quad (7)$$

where $G = \frac{1}{4\sigma^2} (\mathbf{n}_1 \mathbf{n}_2^\top + \mathbf{n}_2 \mathbf{n}_1^\top)$ is symmetric. Since $\tilde{I}(\mathbf{x})$ represents an idealized contrast pattern rather than an actual brightness value, we obtain the image intensity via a linear camera response applied to $\tilde{I}(\mathbf{x})$ [29]:

$$I(\mathbf{x}) = I_0 + \alpha \tilde{I}(\mathbf{x}), \quad (8)$$

where α denotes the camera gain and I_0 offset.

Having established the local intensity model of the checkerboard corner, we now analyze ER on checkerboard

corner by substituting Eq. (8) into the ER formulation in Eq. (3). Since $L(\mathbf{x}) = \log I(\mathbf{x})$, we have

$$\nabla L(\mathbf{x}) \approx \frac{\alpha}{I_0 + \alpha \tilde{I}(\mathbf{x})} G \mathbf{x}. \quad (9)$$

Since $\tilde{I}(\mathbf{x}) \approx 0$ around the corner, the term $\alpha/(I_0 + \alpha \tilde{I}(\mathbf{x}))$ can be approximated by α/I_0 and treated as constant. Therefore, Eq. (3) can be rewritten as

$$\begin{aligned} R_e(\mathbf{x}) &\approx \frac{1}{C} |\nabla L(\mathbf{x}) \cdot \mathbf{v}| \\ &\approx \gamma |(G \mathbf{x}) \cdot \mathbf{v}| = \gamma |(G \mathbf{v})^\top \mathbf{x}|, \end{aligned} \quad (10)$$

where $\gamma = \alpha/C I_0$ is a constant factor. Since G cannot be zero unless $\mathbf{n}_1$ and $\mathbf{n}_2$ are parallel, ER becomes zero at the exact corner $\mathbf{x} = \mathbf{0}$ or in the degenerate case $\mathbf{v} = \mathbf{0}$. Thus, from Eq. (10), we find that

$$\mathbf{x} = \mathbf{0} \implies R_e(\mathbf{0}) = 0. \quad (11)$$

This result holds locally near a checkerboard corner, where the image gradients are modeled as a pair of intersecting edges. In contrast, the intensity model of the checkerboard edge can be modeled as

$$\tilde{I}(\mathbf{x}) \approx \frac{1}{2\sigma} (\mathbf{n}^\top \mathbf{x}), \quad \nabla \tilde{I}(\mathbf{x}) \approx \frac{\mathbf{n}}{2\sigma}. \quad (12)$$

Here, the checkerboard edge corresponds to the set of points $\mathbf{x}$ that satisfy $\mathbf{n}^\top \mathbf{x} = 0$. Then, we can formulate ER on the checkerboard edge as follows:

$$R_e(\mathbf{x}) \approx \kappa |\mathbf{n} \cdot \mathbf{v}|, \quad \text{where } \kappa = \gamma/2\sigma. \quad (13)$$

This expression indicates that events are reliably generated along the edge as long as the motion vector is not parallel to the edge direction.

In conclusion, under the first-order assumption on local models, no events are generated at the checkerboard corner itself, whereas along an edge, events occur whenever $\mathbf{n} \cdot \mathbf{v} \neq 0$. Consequently, ER on edges is strictly higher than at the corner. This contrast between edges and corners forms the basis of our detection pipeline: edge regions provide reliable cues for initialization, while the absence of events at the corner provides a principled criterion for refinement toward the true corner location.

3.3. Event Representation

Event cameras capture visual information asynchronously. This asynchronous stream forms spatio-temporal trajectories based on relative scene motion or ego-motion [11, 12]. Thus, simple accumulation of these events on a static image plane causes motion blur. To mitigate it, various event representations have been proposed [11, 12, 15, 30, 51]. Among them, we adopt IWE [11, 12], which warps events




Figure 3. **Event representation.**

to a reference time under a motion model to enforce geometric alignment and suppress motion blur. In contrast to time-surfaces [30, 51] that encode temporal recency at each pixel, IWE aligns spatio-temporal event trajectories onto a common instant, yielding an edge image consistent with the single-instant assumption of camera calibration [48].

Events Warping. To construct IWE, we first divide the input event stream $\mathcal{E}$ into M coherent windows, $\mathcal{E} = \bigcup_{i=1}^M \mathcal{E}_i$, and process each window $\mathcal{E}_i$ independently. For each window, we warp each event to reference time t_{ref} , using the warping function [11]:

$$e_k = (\mathbf{x}_k, t_k, p_k) \mapsto e'_k = (\mathbf{x}'_k, t_{\text{ref}}, p_k), \quad (14)$$

$$\mathbf{x}'_k = \mathbf{x}_k + (t_{\text{ref}} - t_k) \mathbf{v},$$

where $\mathbf{v} \in \mathbb{R}^2$ denotes the constant image-plane motion vector and $\mathbf{x}'_k$ is the pixel location of the warped event e'_k . This model assumes linear camera motion within each window, which is a valid assumption in calibration scenarios where data can be collected under a controlled environment [48].

Then, we build Image of Warped Events (IWE):

$$H(\mathbf{x}; \mathbf{v}) = \sum_{k=1}^{N_e} p_k \delta(\mathbf{x} - \mathbf{x}'_k), \quad (15)$$

where N_e represents the number of events in each window, each pixel $\mathbf{x}$ accumulates the number of warped events whose locations $\mathbf{x}'_k$ spatially coincide with it, and $\delta(\cdot)$ indicates the Dirac delta function. To estimate the optimal motion vector $\mathbf{v}^*$, we maximize the variance of IWE:

$$\mathbf{v}^* = \arg \max_{\mathbf{v} \in \mathbb{R}^2} \text{Var}(H(\mathbf{x}; \mathbf{v})). \quad (16)$$

This objective encourages sharper and more coherent event alignment [11, 12].

Using the estimated motion vector $\mathbf{v}^*$, we can construct IWE and polarity-aware IWE (pIWE) for each window. IWE $H(\mathbf{x}; \mathbf{v}^*)$ aggregates events regardless of their polarity by setting $p_k = 1$ for all events. It stably captures the landscape of event density without being affected by polarity cancellation. In contrast, pIWE $H_p(\mathbf{x}; \mathbf{v}^*)$ holds the polarity p_k of each event. Using pIWE is effective for identifying black-white alternation of checkerboard patterns, as it preserves polarity contrast. We use these event representations in the subsequent corner detection pipeline (see Fig. 3).

We note that, unlike grayscale images from the reconstruction network [39], which estimates absolute scene




Figure 4. **Comparison with conventional methods.** Harris [21] and Shi-Tomasi [42] fail under event data, whereas the proposed method achieves accurate and stable localization.

brightness, IWE and pIWE are motion-compensated event representations that retain event density. Consequently, ER is directly compatible with these representations.

3.4. Corner Detection and Calibration

Here, we propose our calibration framework, including a checkerboard corner detection pipeline built upon our analytical insights from Sec. 3.2. Our framework consists of four steps: (i) patch extraction, (ii) corner initialization, (iii) corner refinement, and (iv) camera parameters optimization. First, we extract local patches around checkerboard corners to isolate candidate regions. Second, for each patch, we initialize corner positions by identifying local patterns consistent with the geometry of the checkerboard corner. Third, we refine the initial corners to sub-pixel accuracy by optimization. Finally, we aggregate the refined corners and use them for calibration.

Patch Extraction. To pinpoint the refined corner $\mathbf{c}^*$ precisely, we extract small patches to capture local pattern near each corner. While our overall calibration pipeline primarily utilizes IWE, we leverage polarity contrast of pIWE during the patch extraction to detect initial corners.

To extract each patch in pIWE, we adopt circular boundaries inspired by Bok *et al.* [2]. Since a checkerboard corner lies at the intersection of black and white regions, we select pixels whose circular boundaries contain both positive and negative polarities. Then, the selected pixels close to each other are merged into the same clusters. Each cluster is enclosed by its smallest bounding rectangle within pIWE. Since IWE shares the same pixel coordinates with pIWE, we can extract the corresponding candidate corner regions from IWE. We denote such candidate patches as $\mathcal{P}$.

At first glance, it seems to be possible to apply conventional detection algorithms [21, 42] to IWE. However, even with IWE, corner detection remains highly unreliable on event data (see Fig. 4). As analyzed in Eq. (11), this is because corner regions barely trigger any events.

Corner Initialization. We guess the initial corner $\mathbf{c}_0$ within the single patch $\mathbf{P} \in \mathcal{P}$. Although invalid patches might still exist in the candidate patches $\mathcal{P}$, we can filter out invalid patches during this process.




Figure 5. **Pixel propagation from seed pixels.** Each red point means seed pixel $s \in \mathcal{S}$.




Figure 6. **Corner initialization procedure.**

As we analyzed in Sec. 3.2, events rarely occur at corners, and thus cannot be directly used for corner localization. To handle this issue, we first leverage edge-based cues to estimate the initial corner $\mathbf{c}_0$. As each pixel value in IWE represents the event density, pixels with high values correspond to edges where distinct motion occurs. Leveraging this property, we can define seed pixels $\mathcal{S}$ as pixels that have higher values than their eight neighboring pixels. Seed pixels $\mathcal{S}$ are highly likely to belong to certain edges, so we detect them from the boundary regions within patch $\mathbf{P}$ (see Fig. 5). The seed pixels $\mathcal{S}$ are defined as follows:

$$\mathcal{S} = \{\mathbf{x} \in \mathbf{P} \mid H(\mathbf{x}) > H(\mathbf{y}), \forall \mathbf{y} \in \mathcal{N}_8(\mathbf{x})\}, \quad (17)$$

where $\mathbf{x}$ is a pixel in the patch $\mathbf{P}$ and $\mathcal{N}_8(\mathbf{x})$ represents eight neighbors of $\mathbf{x}$. To initialize precise corner points, we propagate from each seed pixel $s \in \mathcal{S}$ toward the second-largest neighbor in $\mathcal{N}_8(s)$ step by step (see Fig. 5). After each seed pixel s propagates one step toward its neighboring pixel, it merges with the propagated pixel, forming a cluster. Let $N(s)$ denote the number of non-empty clusters after merging at step $s \in \mathbb{R}$. Each patch generally maintains $N(s) = 4$, corresponding to the four edges of the checkerboard pattern. When $N(s) \neq 4$, it indicates either a single-edge response or a misdetected seed, and such cases are filtered out. As propagation continues, clusters gradually merge, decreasing $N(s)$. We therefore define s^* as the largest propagation step for which $N(s) = 4$:

$$s^* = \max \{s \in \mathbb{R} \mid N(s) = 4\}. \quad (18)$$

At step s^* , we order the four clusters counterclockwise by label $l \in \{1, 2, 3, 4\}$. We then define two lines based on index parity: $\mathbf{l}_{\text{odd}}$ is the line fitted to all pixels belonging to clusters with an odd label, and $\mathbf{l}_{\text{even}}$ is the line fitted to all pixels belonging to clusters with an even label. Then, we determine the corner $\mathbf{c}_0$ as the intersection $\mathbf{c}_0 = \mathbf{l}_{\text{odd}} \times \mathbf{l}_{\text{even}}$ (see Fig. 6). We can filter out the invalid patch if the two lines are nearly parallel or if obtained $\mathbf{c}_0$ not in $\mathbf{P}$. This approach enables us to initialize a reliable corner $\mathbf{c}_0$ when few or no events occur around the actual corner. The filtered cases are shown in the supplementary material.




Figure 7. **Visualization of corner refinement on the IWE surface.** Its height represents the count of accumulated events. We show the optimization of an initial corner $\mathbf{c}_0$ to its refined corner $\mathbf{c}^*$ along the gradient path.

Corner Refinement. To obtain the sub-pixel accurate refined corner $\mathbf{c}^*$, we refine the initial corner $\mathbf{c}_0$. Based on our analysis in Eq. (11), we assume that the ideal corner corresponds to the minimum value in IWE. Thus, we apply gradient descent to $\mathbf{c}_0$ to find this local minimum (see Fig. 7). During this process, non-corner regions in IWE (*e.g.*, checkerboard squares) may also yield near-zero responses, potentially leading to incorrect convergence. To mitigate this effect, we constrain the optimization to keep corner $\mathbf{c}^*$ within the intersection of two narrow bands around the fitted lines $\mathbf{l}_{\text{odd}}$ and $\mathbf{l}_{\text{even}}$, as follows:

$$\mathbf{c}^* = \arg \min_{\mathbf{c}} H(\mathbf{c}) \quad \text{s.t.} \quad |\mathbf{c}^\top \mathbf{l}_{\text{odd}}| \leq \epsilon, |\mathbf{c}^\top \mathbf{l}_{\text{even}}| \leq \epsilon, \quad (19)$$

where $\mathbf{c}$ denotes a sub-pixel coordinate within the local neighborhood around $\mathbf{c}_0$, and ϵ is a small tolerance.

This refinement process encourages the initial corners to move toward the minimum of the event density within the constrained region, achieving sub-pixel accuracy. Through this process, we obtain geometrically consistent corner positions that align with our theoretical analysis.

Camera Parameter Optimization. Finally, we perform a batch optimization to estimate the camera parameters. Each of the event windows yields a set of refined 2D corners $\mathcal{C}_i^{2D}$, where $i \in \{1, \dots, M\}$. Each element in $\mathcal{C}_i^{2D}$ corresponds to a refined corner $\mathbf{c}^*$ obtained in the previous process. We also denote the 3D corner coordinates of the calibration board as $\mathcal{C}^{3D}$, which are predefined and fixed during optimization. The refined corner sets from all event windows are jointly used with the predefined 3D corner coordinates to optimize the intrinsic matrix $\mathbf{K}$ and distortion coefficients $\mathbf{d}$ by minimizing the reprojection error using `calibrateCamera()` in OpenCV [4]:

$$\mathbf{K}^*, \mathbf{d}^* = \arg \min_{\mathbf{K}, \mathbf{d}, \mathbf{T}_i} \sum_{i=1}^M \|\mathcal{C}_i^{2D} - \pi(\mathbf{K}, \mathbf{d}, \mathbf{T}_i, \mathcal{C}^{3D})\|^2, \quad (20)$$

where $\mathbf{T}_i$ denotes the extrinsic parameters (camera pose) for the i -th event window, and $\pi(\cdot)$ represents the pinhole projection model including distortion.




Figure 8. AprilTag detection procedure.

3.5. AprilTag Detection

By using the polarity contrast of pIWE (Sec. 3.4), our method can be readily extended to AprilTag detection [36, 46]. AprilTag encodes unique IDs within its patterns, enabling reliable identification even under partial visibility or occlusion. Thus, we demonstrate an extension built upon the proposed pattern detection to ID-based fiducial (see Fig. 8). AprilTag detection consists of three steps: (i) quad fitting, (ii) binarization, and (iii) tag decoding. Under the same assumptions in Sec. 3.1, we capture AprilTag patterns using an event camera. By applying our pattern detection method, we can reliably obtain refined corners $\mathcal{C}_i^{2D}$ corresponding to each divided window. For simplicity, we omit the subscript i in the following description.

Quad Fitting. Given the refined corners $\mathcal{C}^{2D}$, we enumerate four-point combinations to obtain quadrilateral candidates $\mathcal{Q}$ (see Fig. 8 (a)). Thus, each candidate $\mathcal{Q}$ is defined as a subset of four corners:

$$\mathcal{Q} \subset \mathcal{C}^{2D}, \quad |\mathcal{Q}| = 4,$$

where $|\cdot|$ denotes the number of elements.

We order each subset $\mathcal{Q}$ clockwise and check whether it satisfies the following geometric constraints to verify its validity as a quadrilateral:

- **Angle range constraint:** All internal angles should remain within a valid range.
- **Minimum area threshold:** The quadrilateral must have sufficient area to ensure geometric stability for detection.
- **Edge length consistency:** The quadrilateral should maintain similar side lengths to prevent distorted shapes.

We then remove overlapping quadrilaterals whose enclosing regions largely overlap using non-maximum suppression (NMS). The remaining quads serve as tag candidates for the subsequent process.

Binarization. In this process, we normalize each valid quadrilateral $\mathbf{q} \in \mathcal{Q}$ into a grid-like binary pattern to easily identify the tag (see Fig. 8 (b)). Based on the four corners of quadrilateral $\mathbf{q}$, we warp the corresponding pIWE region into a normalized square coordinate frame.

From the warped patch, we compute image gradients ∇H_p and project them along motion vector $\mathbf{v}$, which is the same in Sec. 3.3, to obtain a contrast map $g = \nabla H_p \cdot \mathbf{v}$. This operation emphasizes polarity-aligned structural edges of the tag. We then apply ternary thresholding to g to produce $\{-1, 0, +1\}$ values that represent negative edges, neu-

tral regions, and positive edges. Next, we divide the ternarized patch into an $N \times N$ grid and compute the mean value of each cell to construct a coarse binary structure that encodes the tag pattern. Finally, we binarize the cell using Otsu’s adaptive thresholding [37].

Tag Decoding. To assign the tag ID to the given binary pattern, we pack the binary pattern into an integer codeword and compare it with the predefined AprilTag codebook. We achieve rotation invariance by evaluating the Hamming distance across four possible orientations and select the best match when the minimum Hamming distance falls below a predefined threshold [36]. By doing so, we can obtain the unique tag ID and orientation corresponding to given quadrilateral $\mathbf{q}$. Examples of various matching results are provided in the supplementary material.

4. Experiments

4.1. Experimental Setup

Our framework is evaluated on real-world data from a DAVIS346 event camera (346×260 resolution), which provides both asynchronous events and grayscale frames. During acquisition, we move the camera so that all checkerboard edges appear in the view, while keeping the motion approximately linear. Our dataset contains three checkerboard patterns: 5×6 (33 mm), 6×8 (35 mm), and 7×10 (28 mm). In addition to our self-collected data, we also incorporate public data from Prophesee Gen 3 and Samsung Gen 3 event cameras [35]. However, since these public data lack checkerboard information, we use them only for qualitative corner detection evaluation.

To the best of our knowledge, existing event camera calibration methods that support checkerboards rely only on image reconstruction [35, 39]. Therefore, we compare our method against two baselines: (1) E2Calib [35], which uses E2VID [39] for grayscale image reconstruction, and (2) calibration using DAVIS grayscale frames. In our real-world evaluation, we exclude circle-grid calibration methods [6, 23, 41] as they require a different calibration target, and using different targets inevitably leads to different camera–pattern relative poses. Since intrinsic estimation is sensitive to the pose distribution, performance differences would be confounded by target-induced pose variations rather than the calibration method itself [38]. Therefore, achieving a fair comparison across different calibration targets is difficult in real-world settings, and we additionally conduct synthetic experiments, including comparisons with circle-grid calibration [6, 41].

4.2. Quantitative Comparison

Our method demonstrates higher repeatability and lower sensitivity than the comparison methods. As shown in Table 1, it consistently yields smaller standard deviations

Table 1. **Calibration results (mean $\pm$ std) on checkerboard patterns.** Calibration RMSE denotes the root-mean-square reprojection error (in pixels) evaluated using the mean estimated parameters. Frame-based calibration is used as a pseudo ground-truth reference.

Board	Method	Intrinsic Parameters				Distortion Coefficients				Calibration RMSE
		f_x	f_y	c_x	c_y	$k_1 (\times 10^{-1})$	$k_2 (\times 10^{-1})$	$p_1 (\times 10^{-3})$	$p_2 (\times 10^{-3})$	
5 $\times$ 6 (B=32)	E2Calib [35]	332.47 $\pm$ 3.00	332.48 $\pm$ 3.01	187.66 $\pm$ 1.36	131.90 $\pm$ 1.20	-2.39 $\pm$ 0.16	-6.43 $\pm$ 1.24	-3.25 $\pm$ 0.86	-2.93 $\pm$ 0.86	0.747
	Ours	347.96 $\pm$ 1.60	348.51 $\pm$ 1.62	184.25 $\pm$ 0.68	125.43 $\pm$ 0.57	-3.70 $\pm$ 0.09	1.39 $\pm$ 0.72	-1.35 $\pm$ 0.29	-1.47 $\pm$ 0.48	0.487
	Frame-Based	349.30 $\pm$ 1.71	349.99 $\pm$ 1.73	186.82 $\pm$ 0.50	129.57 $\pm$ 0.32	-3.48 $\pm$ 0.05	1.21 $\pm$ 0.20	-0.42 $\pm$ 0.21	-1.90 $\pm$ 0.40	0.199
6 $\times$ 8 (B=32)	E2Calib [35]	346.25 $\pm$ 2.15	346.95 $\pm$ 2.12	185.94 $\pm$ 1.48	133.43 $\pm$ 1.30	-3.67 $\pm$ 0.16	1.82 $\pm$ 0.88	-2.22 $\pm$ 0.60	-0.10 $\pm$ 0.89	0.647
	Ours	347.60 $\pm$ 1.22	347.99 $\pm$ 1.19	183.07 $\pm$ 1.32	131.90 $\pm$ 1.05	-3.45 $\pm$ 0.05	0.86 $\pm$ 0.36	-2.00 $\pm$ 0.50	0.60 $\pm$ 0.60	0.516
	Frame-Based	349.07 $\pm$ 0.87	349.33 $\pm$ 0.86	187.63 $\pm$ 0.67	130.84 $\pm$ 0.49	-3.55 $\pm$ 0.04	2.11 $\pm$ 0.23	-1.60 $\pm$ 0.20	0.40 $\pm$ 0.30	0.192
7 $\times$ 10 (B=25)	E2Calib [35]	342.87 $\pm$ 2.02	342.89 $\pm$ 1.93	187.02 $\pm$ 1.07	129.14 $\pm$ 0.43	-3.39 $\pm$ 0.05	0.76 $\pm$ 0.24	-0.51 $\pm$ 0.35	-1.24 $\pm$ 0.57	0.559
	Ours	341.60 $\pm$ 0.81	341.80 $\pm$ 0.78	184.31 $\pm$ 0.38	131.79 $\pm$ 0.50	-3.42 $\pm$ 0.05	2.00 $\pm$ 0.01	-1.16 $\pm$ 0.34	-1.93 $\pm$ 0.23	0.503
	Frame-Based	342.12 $\pm$ 0.74	342.31 $\pm$ 0.71	190.18 $\pm$ 0.39	130.39 $\pm$ 0.28	-3.56 $\pm$ 0.02	2.01 $\pm$ 0.05	-1.10 $\pm$ 0.19	-3.15 $\pm$ 0.21	0.153

Table 2. **Corner RMSE (mean $\pm$ std).** RMSE is the Euclidean distance (in pixels) between checkerboard corner locations from each method and the frame-based method.

Board	E2Calib [35]	Ours
5 $\times$ 6	1.3979 $\pm$ 0.8350	0.8504 $\pm$ 0.3203
6 $\times$ 8	1.0377 $\pm$ 0.4879	0.8356 $\pm$ 0.4776
7 $\times$ 10	0.7808 $\pm$ 0.2742	0.7791 $\pm$ 0.2742

for all intrinsic and distortion parameters across the three checkerboard patterns, indicating more stable calibration. For this evaluation, we randomly sample the set of B corresponding images in which all modalities (E2Calib [35], ours, and frame-based) detect checkerboard corners successfully and repeat calibration 50 times for each method.

To verify the accuracy of our method, we evaluate corner localization accuracy by measuring the Euclidean distance between the corner locations projected using the camera parameters estimated by each method and the projected locations of the reference parameters obtained from the DAVIS grayscale frames (see Table 2). It shows that our method achieves markedly lower RMSE compared to E2Calib (*i.e.*, E2VID-based calibration). Additional quantitative results, including evaluations on synthetic data and computational time, are provided in the supplementary material.

4.3. Qualitative Results

As shown in Fig. 9, our method yields geometrically consistent corner detections on both self-collected and public data. The extracted corners remain well localized, providing reliable inputs for subsequent calibration. Beyond checkerboards, we showcase that our approach extends to various fiducial markers such as AprilTag [36] (see Fig. 10). Our method successfully detects AprilTag with various IDs and geometric configurations, and handles partial board observations by correctly indexing the visible subset of tags even when others fall outside the field of view. Additional qualitative results in the supplementary material include detec-




Figure 9. **Detected checkerboard corners.** The left side of the dotted line is self-collected. The right side is public [35].




Figure 10. **Qualitative results on various AprilTag.**

tions across diverse poses and, for AprilTag, corner indexing under partial observations.

5. Conclusion

We presented the first event-based calibration framework that directly detects checkerboard corners without grayscale image reconstruction. Guided by a mathematical analysis revealing the absence of events at corner intersections, our method leverages edge-driven initialization and event density-based refinement to achieve sub-pixel corner accuracy. Extensive experiments demonstrate that our approach achieves accurate calibration precision. By reintroducing checkerboards and extending them to fiducial markers such as AprilTag, this work showcases the feasibility of real-world deployment for event camera calibration.

Limitation. Our method assumes a dominant global motion and estimates a single motion vector in IWE generation stage, which may degrade performance in scenes with spatially varying motion. Incorporating more general motion models (*e.g.*, pixel-wise motion) remains future work.

Acknowledgments. This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.RS-2020-II201336, Artificial Intelligence Graduate School Program (UNIST); No.RS-2022-II220907, Development of AI Bots Collaboration Platform and Self-organizing; No.RS-2022-II220612, Geometric and Physical Commonsense Reasoning based Behavior Intelligence for Embodied AI), and the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.RS-2024-00457065).

References

- [1] Abdulla Ayyad, Mohamad Halwani, Dewald Swart, Rajkumar Muthusamy, Fahad Almaskari, and Yahya Zweiri. Neuromorphic vision based control for the precise positioning of robotic drilling systems. *Robotics and Computer-Integrated Manufacturing*, 79:102419, 2023.
- [2] Yunsu Bok, Hyowon Ha, and In So Kweon. Automated checkerboard detection and indexing using circular boundaries. *Pattern Recognition Letters*, 2016.
- [3] Jean-Yves Bouguet. Camera calibration toolbox for matlab, 2003.
- [4] Gary Bradski and Adrian Kaehler. *Learning OpenCV: Computer Vision with the OpenCV Library*. O'Reilly Media, 2008.
- [5] Bolin Cai, Ami Zi, Jun Yang, Guoliang Li, Yang Zhang, Qiujie Wu, Chenen Tong, Wenxiang Liu, and Xiangcheng Chen. Accurate event camera calibration with fourier transform. *IEEE TIM*, 73:1–12, 2024.
- [6] Shuolong Chen, Xingxing Li, Liu Yuan, and Ziao Liu. ekalibr: Dynamic intrinsic calibration for event cameras from first principles of events. *IEEE RA-L*, 10(7):7094–7101, 2025.
- [7] Arturo De la Escalera and Jose María Armingol. Automatic chessboard detection for intrinsic and extrinsic camera parameter calibration. *Sensors*, 10(3):2027–2044, 2010.
- [8] Davide Falanga, Kevin Kleber, and Davide Scaramuzza. Dynamic obstacle avoidance for quadrotors with event cameras. *Science Robotics*, 5(40):eaaz9712, 2020.
- [9] Mark Fiala. Artag, a fiducial marker system using digital techniques. In *CVPR*, 2005.
- [10] Paul Furgale, Joern Rehder, and Roland Siegwart. Unified temporal and spatial calibration for multi-sensor systems. In *IEEE IROS*, 2013.
- [11] 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 *CVPR*, 2018.
- [12] Guillermo Gallego, Mathias Gehrig, and Davide Scaramuzza. Focus is all you need: Loss functions for event-based vision. In *CVPR*, 2019.
- [13] Guillermo Gallego et al. Event-based vision: A survey. *IEEE TPAMI*, 44(1):154–180, 2020.
- [14] Sergio Garrido-Jurado, Rafael Muñoz-Salinas, Francisco José Madrid-Cuevas, and Manuel Jesús Marín-Jiménez. Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition*, 47(6):2280–2292, 2014.
- [15] Luna Gava, Marco Monforte, Massimiliano Iacono, Chiara Bartolozzi, and Arren Glover. Puck: Parallel surface and convolution-kernel tracking for event-based cameras. *arXiv preprint arXiv:2205.07657*, 2022.
- [16] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. EKL: Asynchronous, photometric feature tracking using events and frames. *IJCV*, 128(3):601–618, 2020.
- [17] Mathias Gehrig and Davide Scaramuzza. Recurrent vision transformers for object detection with event cameras. In *CVPR*, 2023.
- [18] Mathias Gehrig, Willem Aarents, Daniel Gehrig, and Davide Scaramuzza. Dsec: A stereo event camera dataset for driving scenarios. *IEEE RA-L*, 6(3):4947–4954, 2021.
- [19] Daniel Gorchard. DVScalibration. <https://github.com/gorchard/DVScalibration.git>, 2025.
- [20] Thomas Gossard, Andreas Ziegler, Levin Kolmar, Jonas Tebbe, and Andreas Zell. ewand: An extrinsic calibration framework for wide baseline frame-based and event-based camera systems. In *IEEE ICRA*, 2024.
- [21] Chris Harris, Mike Stephens, et al. A combined corner and edge detector. pages 10–5244, 1988.
- [22] Janne Heikkilä. Geometric camera calibration using circular control points. *IEEE TPAMI*, 22(10):1066–1077, 2002.
- [23] Kun Huang, Yifu Wang, and Laurent Kneip. Dynamic event camera calibration. In *IEEE IROS*, 2021.
- [24] Kunping Huang, Sen Zhang, Jing Zhang, and Dacheng Tao. Event-based simultaneous localization and mapping: A comprehensive survey, 2024.
- [25] Javier Ibáñez, José M. Alonso, Jorge Sastre, Emilio Defez, and Pedro Alonso-Jordá. Advances in the approximation of the matrix hyperbolic tangent. *Mathematics*, 9(11), 2021.
- [26] Juho Kannala and Sami S Brandt. A generic camera model and calibration method for conventional, wide-angle, and fish-eye lenses. *IEEE TPAMI*, 28(8):1335–1340, 2006.
- [27] Juho Kannala, Janne Heikkilä, and Sami S Brandt. Geometric camera calibration. *Wiley encyclopedia of computer science and engineering*, 13(6):1–20, 2008.
- [28] Jun-Sik Kim, Pierre Gurdjos, and In-So Kweon. Geometric and algebraic constraints of projected concentric circles and their applications to camera calibration. *IEEE TPAMI*, 27(4): 637–642, 2005.
- [29] Lars Krüger and Christian Wöhler. Accurate checkerboard corner localisation for camera calibration. *Pattern Recognition Letters*, 32(10):1428–1435, 2011.
- [30] Xavier Lagorce, Garrick Orchard, Francesco Galluppi, Bertram E. Shi, and Ryad B. Benosman. Hots: A hierarchy of event-based time-surfaces for pattern recognition. *IEEE TPAMI*, 39(7):1346–1359, 2017.
- [31] Jianing Li, Jia Li, Lin Zhu, Xijie Xiang, Tiejun Huang, and Yonghong Tian. Asynchronous spatio-temporal memory network for continuous event-based object detection. *IEEE TIP*, 31:2975–2987, 2022.

- [32] Shijie Lin, Yinqiang Zhang, Lei Yu, Bin Zhou, Xiaowei Luo, and Jia Pan. Autofocus for event cameras. In *CVPR*, 2022.
- [33] Nathan Matsuda, Oliver Cossairt, and Mohit Gupta. MC3D: Motion contrast 3D scanning. In *ICCP*, 2015.
- [34] Nico Messikommer, Carter Fang, Mathias Gehrig, and Davide Scaramuzza. Data-driven feature tracking for event cameras. In *CVPR*, 2023.
- [35] Manasi Muglikar, Mathias Gehrig, Daniel Gehrig, and Davide Scaramuzza. How to calibrate your event camera. In *CVPRW*, 2021.
- [36] Edwin Olson. Apriltag: A robust and flexible visual fiducial system. In *IEEE ICRA*, 2011.
- [37] Nobuyuki Otsu et al. A threshold selection method from gray-level histograms. *Automatica*, 11(285-296):23–27.
- [38] Lala Shakti Swarup Ray, Bo Zhou, Lars Krupp, Sungho Suh, and Paul Lukowicz. A synthetic benchmarking pipeline to compare camera calibration algorithms. In *ICPR*, 2024.
- [39] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. High speed and high dynamic range video with an event camera. *IEEE TPAMI*, 43(6):1964–1980, 2021.
- [40] Robotics and Perception Group. rpg_dvs_ros. https://github.com/uzh-rpg/rpg_dvs_ros, 2025. Accessed: Jan. 3, 2025.
- [41] Mohammed Salah, Abdulla Ayyad, Muhammad Humais, Daniel Gehrig, Abdelqader Abusafieh, Lakmal Seneviratne, Davide Scaramuzza, and Yahya Zweiri. E-calib: A fast, robust, and accurate calibration toolbox for event cameras. *IEEE TIP*, 33:3977–3990, 2024.
- [42] Jianbo Shi and Tomasi. Good features to track. In *CVPRW*, pages 593–600, 1994.
- [43] Chaehyeon Song, Dongjae Lee, Jongwoo Lim, and Ayoung Kim. Camera calibration via circular patterns: A comprehensive framework with measurement uncertainty and unbiased projection model. *IEEE TPAMI*, 2025.
- [44] Sihao Sun, Giovanni Cioffi, Coen De Visser, and Davide Scaramuzza. Autonomous quadrotor flight despite rotor failure with onboard vision sensors: Frames vs. events. *IEEE RA-L*, 6(2):580–587, 2021.
- [45] Varun Sundar, Matthew Dutson, Andrei Ardelean, Claudio Bruschini, Edoardo Charbon, and Mohit Gupta. Generalized event cameras. In *CVPR*, 2024.
- [46] John Wang and Edwin Olson. Apriltag 2: Efficient and robust fiducial detection. In *IEEE IROS*, 2016.
- [47] Yongqing Wang, Shiyu He, Yufan Fei, and Xingjian Liu. Motion-error-free calibration of event camera systems using a flashing target. *Optics Express*, 32(15):26833–26845, 2024.
- [48] Zhengyou Zhang. A flexible new technique for camera calibration. *IEEE TPAMI*, 22(11):1330–1334, 2002.
- [49] Zelin Zhang, Anthony J. Yezzi, and Guillermo Gallego. Formulating event-based image reconstruction as a linear inverse problem with deep regularization using optical flow. *IEEE TPAMI*, 45(7):8372–8389, 2023.
- [50] Yajing Zheng, Zhaofei Yu, Song Wang, and Tiejun Huang. Spike-based motion estimation for object tracking through bio-inspired unsupervised learning. *IEEE TIP*, 32:335–349, 2022.
- [51] Shifan Zhu, Zhipeng Tang, Michael Yang, Erik Learned-Miller, and Donghyun Kim. Event camera-based visual odometry for dynamic motion tracking of a legged robot using adaptive time surface. In *IEEE IROS*, 2023.