

Density Invariant Contrast Maximization for Neuromorphic Earth Observations

Sami Arja^{*1}, Alexandre Marcireau¹, Richard L. Balthazor²
 Matthew G. McHarg², Saeed Afshar¹ and Gregory Cohen¹

¹Western Sydney University

²United States Air Force Academy

^{*}s.elarja@westernsydney.edu.au

Abstract

Contrast maximization (CMax) techniques are widely used in event-based vision systems to estimate the motion parameters of the camera and generate high-contrast images. However, these techniques are noise-intolerance and suffer from the multiple extrema problem which arises when the scene contains more noisy events than structure, causing the contrast to be higher at multiple locations. This makes the task of estimating the camera motion extremely challenging, which is a problem for neuromorphic earth observation, because, without a proper estimation of the motion parameters, it is not possible to generate a map with high contrast, causing important details to be lost. Similar methods that use CMax addressed this problem by changing or augmenting the objective function to enable it to converge to the correct motion parameters. Our proposed solution overcomes the multiple extrema and noise-intolerance problems by correcting the warped event before calculating the contrast and offers the following advantages: it does not depend on the event data, it does not require a prior about the camera motion and keeps the rest of the CMax pipeline unchanged. This is to ensure that the contrast is only high around the correct motion parameters. Our approach enables the creation of better motion-compensated maps through an analytical compensation technique using a novel dataset from the International Space Station (ISS). Code is available at https://github.com/neuromorphicsystems/event_warping

1. Introduction

Neuromorphic event cameras [6,24] are biology-inspired optical sensors that offer high-speed, high-dynamic range, and low data rate operation, making them extremely well suited for use in the space environment. The sensors are asynchronous and have in-pixel circuitry to produce high-temporal contrast change events only in response to changes

Event Camera Observing Earth from the ISS

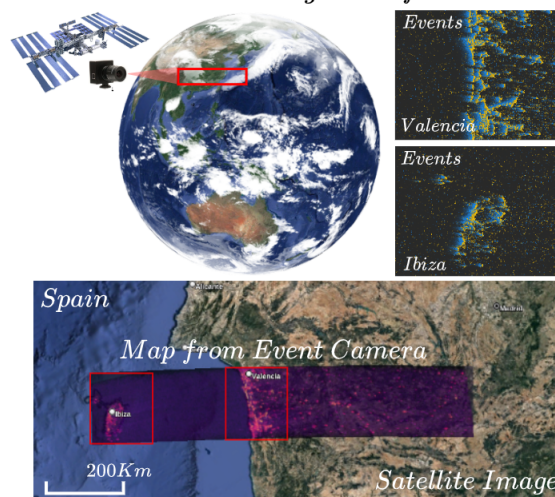


Figure 1. Neuromorphic Earth Observations application. The Falcon Neuro project placed two event cameras on the ISS in 2021. These sensors have been used for earth observation and have captured a dataset of a variety of different locations in 2022. To produce panoramic map images from the event-based output of these sensors, a contrast maximization algorithm, such as CMax [8], is needed to compensate for the motion of the ISS. In this paper, we proposed an analytical approach that creates high-contrast panoramic images using CMax [8] by making the algorithm invariant to the density of events by leveraging the physical properties of the events and their geometries. This was performed by introducing a volumetric-based piecewise correction function that adjusts the warped image to prevent the loss landscape from forming multiple extrema.

in the visual scene. The change events can be represented as event $e = (\mathbf{u}, p, t)$, where $\mathbf{u} = (x, y)$ is the pixel coordinates of the event occurrence, p is the polarity of the contrast change, corresponding to whether the change in the brightness is increasing or decreasing, t is the timestamp of the

change in the order of μs .

The asynchronous nature of the events provide numerous advantages over traditional vision sensors, such as superior dynamic range, low latency, high temporal resolution and significantly lower power consumption. The high temporal resolution removes the effects of motion blur, and the change detection suppresses redundant information at the pixel level. These features make the sensors well suited to tackling challenging machine vision tasks such as recognition [1, 3, 23, 36, 52], tracking [2, 12, 14, 22, 33, 34, 41, 55, 63], SLAM [11, 19, 21, 44, 56, 58], motion estimation [7–9, 25, 26, 30, 35, 37, 38, 40, 47, 48, 50, 54, 59], and space domain awareness and space imaging [1, 4, 5, 18, 42]. The lack of a frame-based output requires the development of new algorithms and systems, and can allow for modes of operation and sensing not possible with conventional imaging sensors.

The recent advances in event-based algorithms and the wide availability of high-resolution event cameras [6] have led to wide adoption in numerous real-world research applications such as in Autonomous Underwater Vehicles (AUV) [60], ground-based mobile telescopes [4, 5], Unmanned Aerial Vehicles (UAV) [32, 43, 57], ground robots and vehicles [17, 28, 31], and even in space onboard the ISS [29].

There is increasing interest in using event cameras in the space environment and further afield. These include the investigation of their use in future lunar spacecraft landing tasks [30, 51] and for underground exploration using the Mars Ingenuity Helicopter [27]. These projects used either simulated events from video sequences or earth-based environments that resemble the Martian and the Lunar surfaces.

In this paper, we investigate neuromorphic earth observation using event cameras mounted on the ISS. These sensors were installed in 2021 as part of a collaboration between the United States Air Force Academy and Western Sydney University through the Falcon Neuro project [29]. The Falcon Neuro payload contains two identical neuromorphic sensors, with one pointed in the RAM direction, and the other pointed in the NADIR direction. This work focuses on techniques to process data from the NADIR camera to produce visual maps of the earth through an analytical solution to the original CMax framework.

1.1. Motivation

A state-of-the-art approach for motion estimation was first introduced by [8], known as CMax. It works by estimating the camera's relative motion vector, $\theta = [v_x, v_y]$, over a time window $\delta = t_i - t_{ref}$ to align events with the edges or objects that generated them. This involves adjusting the pixel coordinates of individual events to eliminate the motion induced by the sensor or object to create sharp images. Specifically, each event $e_i = (\mathbf{u}_i, p_i, t_i)$ is warped by a shear transformation based on a motion candidate θ

$$\mathbf{u}'_i = \begin{bmatrix} x'_i \\ y'_i \end{bmatrix} = \begin{bmatrix} x_i \\ y_i \end{bmatrix} - \begin{bmatrix} v_x \\ v_y \end{bmatrix} * \delta \quad (1)$$

This works by reversing the motion θ between t_i and the beginning of δ and changes the spatial location of the \mathbf{u}_i . The new events are then accumulated into an image H or also called Image of Warped Events (IWE) as in (2). Each pixel in (2) sums the values of the warped events \mathbf{u}_i that fall within it.

$$H(\mathbf{u}'; \theta) \doteq \sum_{i=1}^N b_k \delta(\mathbf{u} - \mathbf{u}'_i), \quad (2)$$

where b_k is the number of events along the trajectories as detailed in [8]. The contrast of H is then calculated as a function of θ

$$C(\theta) \doteq \frac{1}{N} \sum_{j=1}^N (H(\mathbf{u}'_j; \theta) - \mu(\theta))^2, \quad (3)$$

where N is the number of pixels in H , and $\mu(\theta)$ is the mean of the pixel intensity of H . The strategy is to find the correct θ that maximises the objective function $C(\theta)$, in this case, a higher $C(\theta)$ indicates higher events alignment and a sharp motion-corrected image H .

Despite the recent successes of the CMax algorithm in various applications, the algorithm suffers from a few fundamental weaknesses particularly, in space-based earth mapping applications caused by the increased density of events resulting from the continuous movement of the camera in the orbit. Below we discuss these limitations and use them to motivate the method proposed in this paper.

This work makes use of the event camera pointed directly downward (the NADIR direction) as shown in Figure 1. The ISS moves with a consistent speed of 17,900 mph (i.e. 8 km/s) [10] and the movement can be modelled as consisting primarily of translation, except during docking procedures. The number of events recorded by the camera is predominantly influenced by the texture of the surface of the earth and the variations in lighting conditions over the day or night side of the earth.

The high texture of the lower atmosphere (e.g. clouds) combined with pixel noise caused by the camera setting and circuit mismatch, reduces the Signal-to-Noise Ratio (SNR). A low SNR scene results in fewer structures captured by the camera, affecting the accuracy of the motion estimation θ . When the density of events in (2) increases, this also leads to higher variance in (3) around the wrong θ . As a consequence, higher contrast does not always imply better events alignment, as illustrated in Figure 2(a), where multiple extrema become visible. This noise intolerance issue was initially identified by [54] but was not further investigated.

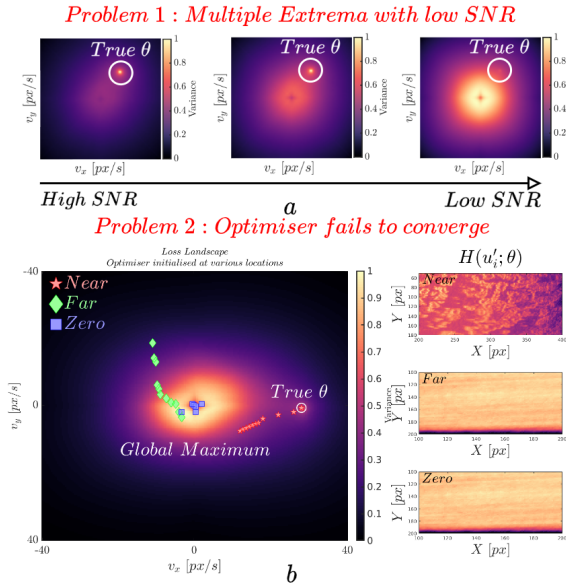


Figure 2. Illustration of CMax problems in low SNR scenes. **Problem 1:** It is trivial to estimate the true motion θ in high SNR scenes, however, as the SNR decreases the contrast around the wrong θ increases significantly, overtaking the true value. **Problem 2:** Optimisation algorithm converges only when it is initialised near the true θ and fails everywhere else, this is a problem because a robust CMax should converge to the correct motion θ without any prior.

Recent studies have shown that maximization (3) can be done either by using a conjugate gradient [8, 9] approach, or a branch-and-bound [25] method. The former requires a good initialisation of θ to converge to the correct local minima, as summarised in Figure 2(b), and the latter is a global optimisation and search method better suited for rotation motion estimation. In addition, recent improvements have included refining the objective function to better suit the targeted settings [7, 54]. However, this makes the objective function and optimisation application-dependent and may increase complexity.

Guided by these observations, we propose a new approach that corrects (2) by only considering the motion and geometry of the warped events, which modifies the landscape of the objective function in (3) automatically. This modification enables us to keep the optimisation algorithm and objective function unchanged in a way invariant to the nature of the input data.

1.2. Contribution

We present a novel approach that enables space-based earth mapping using the CMax algorithm. Our approach makes the following assumptions: the speed of the camera is constant, the time window δ is the entire event stream

and no motion prior is given to the algorithm. Our method not only provides an optimal solution to this specific application, but it also addresses several fundamental problems, including determining the appropriate objective function to use [7] and deciding how many events to process at once [54]. We demonstrate that the variance can serve as an optimal objective function without any modifications, and all events can be utilised in the case of translation motion. While our method does not address other types of motion, such as rotation and zooming, this eliminates the need to test multiple objective functions and employ several batch sizes to estimate the optimal motion parameters.

Our method relies solely on the overall geometry of the warped events and it does not depend on noise density or the size of the time window. It enables CMax to consistently produce high-contrast outputs to equation (3) around the correct θ , even in cases where noise dominates the primary structure of the scene. This is accomplished by increasing the value of specific pixels of the Image of Warped Events – namely, pixels that correspond to parts of the scene in front of which the sensor spent less time. Our approach recovers a single solution for θ . This significantly increases the rate of convergence to the correct solution using a simple optimisation search method such as Nelder-Mead, thus overcoming the problem of multiple extrema.

We model the noise events of noise-rich scenes as a uniform rectangular cuboid in the 3-dimensional space $\{x, y, t\}$. This allows us to analytically calculate the variance of the accumulated image, by shearing the cuboid and integrating the function that describes its height after shearing.

Our method was evaluated on a diverse dataset captured from the ISS. The results were assessed both qualitatively and quantitatively using Root Mean Square (RMS) error and the Rate of Convergence (RoC) metrics for evaluation – see Section 3.1 and 3.2.

1.3. Related Work

The CMax framework was introduced by [8, 9] and its performance was further investigated in [7, 54]. Different alignment methods have been proposed to leverage the benefit of this algorithm [15, 35, 47] for various tasks such as rotational motion estimation [9, 19, 20, 46], optical flow [16, 49, 64], 3D reconstruction [13, 45], depth estimation [8], motion segmentation [39, 53, 62], and intensity reconstruction [61]. To estimate the motion, existing works either rely on local [7, 8, 53, 54], or global optimisation [25, 40] to facilitate the convergence to the correct motion parameters.

[50] is the work most closely related to ours. Their method augments the objective function and applies corrections derived from mathematical models. However, [50] tackles the problem of event collapse, which occurs when objects are not moving parallel to the sensor plane. By con-

trast, we only focus on translational motion and present a method that corrects noise-induced variance. This problem only becomes visible with long time windows and high noise levels.

2. Our Approach

Our method operates on the Image of Warped Events 2, before calculating the image contrast 3. It applies a multiplicative correction that solely depends on the motion candidate θ and the width and height of the sensor. We first describe the method for one-dimensional sensors 2.1. We then expand it to two-dimensional sensors 2.2.

2.1. One-dimensional case

Let us consider an event sensor with a single row of pixels that generates random, uniformly distributed noise events with overall rate ρ (in events per second). These events can be seen as a dense point cloud in two-dimensional space $\{x, t\}$. We approximate this point cloud with a "solid" rectangle in $\{x, t\}$ that spans the width of the sensor. Under this approximation, warping the events is equivalent to shearing the rectangle and results in a parallelogram in $\{x, t\}$. The transformation shifts the top of the rectangle by $-v\delta$, where v is the candidate speed and δ is the considered time window.

We denote by f the *Line of Warped Events* (the one-dimensional equivalent of the Image of Warped Events). Its values $f(x)$ are given by the height of the parallelogram at x , hence their plot is a trapezoid without a base (3). The "contrast" in the Line of Warped Events can be estimated with the variance of f over the interval $[0, w - v\delta]$ (for $v \leq 0$), denoted var_f . Importantly, we are not calculating the variance of a random variable but simply considering a continuous extension of the formula for the variance of a collection of samples. This is similar to the difference between the mean of a random variable and the mean of a function.

$$\text{var}_f(v) = \frac{1}{w - v\delta} \int_0^{w-v\delta} (f(x) - \bar{f})^2 dx \quad \text{if } v \leq 0 \quad (4)$$

\bar{f} is the mean of f over the interval $[0, w - v\delta]$.

$$\bar{f}(v) = \frac{1}{w - v\delta} \int_0^{w-v\delta} f(x) dx \quad \text{if } v \leq 0 \quad (5)$$

For $v \geq 0$, one must consider the interval $[-v\delta, w]$ and divide by $w + v\delta$.

var_f is zero if and only if f is constant. That is the case for $v = 0$, however all other candidate speeds result in a

non-zero variance. This change in variance, which is created by sheared noise, causes the problem described in figure 2. To show that our simple continuous model is sufficient to describe the observations, we calculate below an explicit expression of var_f as a function of v .

Without loss of generality, we can restrict the problem to $v < 0$ and introduce unit-less variables to simplify the expressions.

- Normalised pixel position $p = \frac{x}{w}$
- Normalised shear $s = \frac{v\delta}{w}$
- Normalised event count $c = \frac{\delta}{\rho}$

$f(p)$ is a piecewise linear function made of three segments, with two slightly different expressions for $s \leq 1$ and $s \geq 1$ (figure 4). These two expressions correspond (respectively) to shears smaller than the sensor width and shears larger than the sensor width.




Figure 4. Warping dense noise is similar to shearing a "solid" rectangle. The operation turns the rectangle into a parallelogram. The height of the parallelogram as a function of the pixel position describes the top of a trapezoid, shown here in red. When s is larger than one (the time window times the speed is larger than the sensor width), the curve's maximum value decreases. The variance of the red curve as a function of the shear, in blue, follows a surprisingly complex pattern, with a maximum at $s = \frac{1}{2}$.

For $s \leq 1$, f is defined on $[0, 1 + s]$ by:

$$f(p) = \begin{cases} \frac{cp}{s} & 0 \leq p \leq s \\ c & s \leq p \leq 1 \\ c(1 - \frac{p-1}{s}) & 1 \leq p \leq 1+s \end{cases} \quad (6)$$

Applying the formulas for the mean and variance give for $s \leq 1$:

$$\bar{f} = \frac{c}{s+1} \quad \text{and} \quad \text{var}_f(s) = c^2 \frac{s \cdot (2-s)}{3(s+1)^2} \quad (7)$$

The equations for $s \geq 1$ are given in the supplementary materials (Section II). Plotting var_f yields a figure that is very similar to a section of the velocity landscape obtained from real event sensor data (5). The function var_f admits a maximum at $s = \frac{1}{2}$, when the velocity multiplied by the time window equals half the sensor width.

We want var_f to be zero, since f represents the Line of Warped Events for uniform noise. We thus introduce α , a

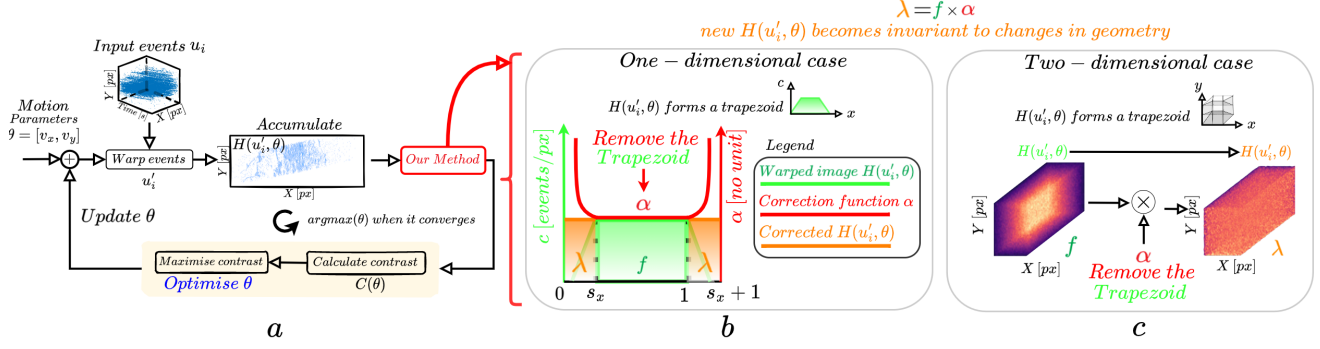


Figure 3. A detailed overview of the proposed approach. **a:** The process of CMax including our method. **b:** The One-dimensional case shows how the trapezoid is formed after accumulating the warped event with the correction function α that removes it. **c:** The Two-dimensional case that take into account both pixel dimensions and removes the trapezoid, producing a new warped image that is invariant to geometry and capable of removing the multiple extrema from the loss landscape.

multiplicative correction function for the non-constant segments of f (the slopes of the trapezoid). α is defined on $[0, 1 + s]$ by:

$$\alpha(p) = \begin{cases} \frac{s}{p} & 0 \leq p \leq s \\ 1 & s \leq p \leq 1 \\ \frac{s}{s+1-p} & 1 \leq p \leq 1 + s \end{cases} \quad (8)$$

Since α is a multiplicative correction rather than an additive one, it does not depend on c (and, by extension, ρ). In other words, the correction can be applied with no prior knowledge of the noise density. The corrected counterpart of f , denoted λ , is defined by:

$$\lambda(p) = \alpha(p) \cdot f(p) \quad (9)$$

By construction, λ is equal to the constant c on $[0, 1 + s]$, and its variance is trivially zero. The figure 5 shows a plot of the variance expression and demonstrates the correction on simulated discrete noise events.

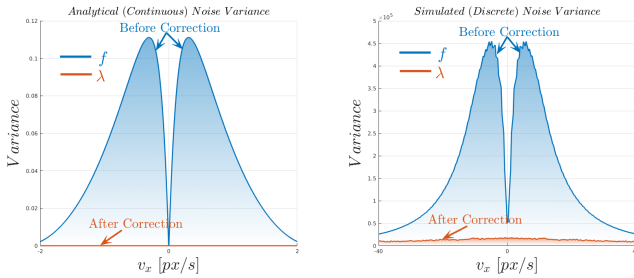


Figure 5. Plots of the variance of f and λ as a function of the candidate speed, using the analytical formula (7) (left) and a discrete simulation (right).

2.2. Two-dimensional case

In this section, we extend our model to two-dimensional sensors. Let us consider an event sensor with $w \times h$ pixels that generates random, uniformly distributed noise events with overall rate ρ (in events per second), a speed candidate $\theta = [v_x, v_y]$, and a time window δ . We model the noise events as a "solid" rectangular cuboid in $\{x, y, t\}$ space.

We define the following normalised variables to express the variance in two dimensions:

- Normalised x pixel position $p_x = \frac{x}{w}$
- Normalised y pixel position $p_y = \frac{y}{h}$
- Normalised shear alongside the x axis $s_x = \frac{v_x \delta}{w}$
- Normalised shear alongside the y axis $s_y = \frac{v_y \delta}{h}$
- Normalised event count $c = \frac{\delta}{\rho}$

Shearing the rectangular cuboid in two directions yields a parallelepiped. Unlike the one-dimensional case, in which the height of the sheared geometry was another well-known figure, the height of the parallelepiped at every point of the $\{x, y\}$ plane forms a "generic" polyhedron with 7 faces without counting the base (figure 7).

Another complication in two dimensions is the shape of the integration domain. While it would be tempting to use a rectangular integration domain, some parts of the scenes are never in the field of view (specifically, two triangular regions of combined area $s_x s_y$, as shown in the Houston recording in figure 8). We take this into account in the integration calculation and divide by $(1 + s_x)(1 + s_y) - s_x s_y$ in the mean and variance formulas.

For $s_x \leq 1$ and $s_y \leq 1$, the height function f is defined by:

$$f(p_x, p_y) = \begin{cases} \frac{cp_y}{s_y} & p_x \leq s_x \wedge p_y \leq \frac{s_y}{s_x p_x} \\ \frac{cp_x}{s_x} & p_x \leq \frac{s_x}{s_y p_y} \wedge p_y \leq s_y \\ \frac{cp_y}{s_y} & s_x \leq p_x < 1 \wedge p_y \leq \cdot \\ \frac{cp_x}{s_x} & p_x \leq s_x \\ c & \wedge s_y \leq p_y \leq 1 \\ c & s_x \leq p_x \leq 1 \\ c & \wedge s_y \leq p_y \leq 1 \\ c \left(\frac{1-p_x}{s_x} + \frac{p_y}{s_y} \right) & 1 \leq p_x \leq 1 + s_x \\ & \wedge (p_x - 1) \frac{s_y}{s_x} \leq p_y \leq s_y \end{cases} \quad (10)$$

Calculating the mean and variance of f yields:

$$\bar{f} = \frac{c}{s_x + s_y + 1} \quad (11)$$

$$\text{var}_f(s_x, s_y) = c^2 \frac{(s_x^2 s_y + s_x s_y^2 - 3s_x s_y + q_x + q_y)}{6(s_x + s_y + 1)^2} \quad (12)$$

where $q_x = -2s_x^2 + 4s_x$ and $q_y = -2s_y^2 + 4s_y$.

Figure 6 shows a plot of the variance as a function of s_x and s_y . Despite its simplicity, our model predicts quite well the "ring" that we observe on real sensor data. Similarly to the one-dimensional case, we introduce a multiplicative correction function α to "flatten" f and ensure that the variance of the corrected height function is zero. For $s_x \leq 1$ and $s_y \leq 1$, α is defined by:

$$\alpha(p_x, p_y) = \begin{cases} \frac{s_y}{p_y} & p_x \leq s_x \wedge p_y \leq \frac{s_y}{s_x p_x} \\ \frac{s_x}{p_x} & p_x \leq \frac{s_x}{s_y p_y} \wedge p_y \leq s_y \\ \frac{s_y}{p_y} & s_x \leq p_x \leq 1 \wedge p_y \leq s_y \\ \frac{s_x}{p_x} & p_x \leq s_x \wedge s_y \leq p_y \leq 1 \\ 1 & s_x \leq p_x \leq 1 \\ & \wedge s_y \leq p_y \leq 1 \\ \frac{s_x s_y}{p_y s_x - s_y p_x + s_y} & 1 \leq p_x \leq 1 + s_x \\ & \wedge (p_x - 1) \frac{s_y}{s_x} \leq p_y \leq s_y \end{cases} \quad (13)$$

The equations for $s_x \geq 1$ and $s_y \geq 1$ are given in the supplementary materials (Section III). To obtain a corrected warped image from real event data, we can directly apply the correction α to the pixels of the accumulated image (2), given only the candidate warp speed $[v_x, v_y]$ and the dimensions of the sensor. This minimises the contribution of uniform noise to the variance in the image and produces a new accumulated frame denoted as $H(\mathbf{u}'_i, \boldsymbol{\theta})_v$, as described in (14). Finally, we can calculate the variance using (3) on the corrected image to solve the problem.




Figure 6. The left graph (a) is a plot of the analytical formula for the variance as a function of the candidate speed. The right graph (b) is a plot of the corrected variance calculated on discrete simulated noise, and is zero as intended.

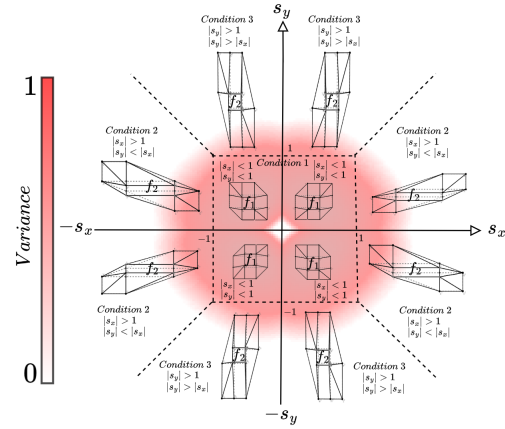


Figure 7. An illustration of the height of the sheared rectangular cuboid (2) as a function of s_x and s_y (black geometric figures) and the corresponding variance (red and white background). The problem exhibits several symmetries and can be solved by considering only two sets of conditions (condition 2 and condition 3 are symmetrical about the axis $y = x$).

$$H(\mathbf{u}'_i, \boldsymbol{\theta})_v = H(\mathbf{u}'_i, \boldsymbol{\theta}) \times \alpha \quad (14)$$

We employed analytical integration techniques based on motion and geometry to ensure that the variance is only high around correct $\boldsymbol{\theta}$. By extending our approach to 2D, we can now apply this corrective technique to real-world data. In the next section, we demonstrate how this method can be used to generate motion-compensated images from data acquired from the ISS.

3. Experiments

To show how our method successfully generalises to real event data, we now apply this technique to data in which

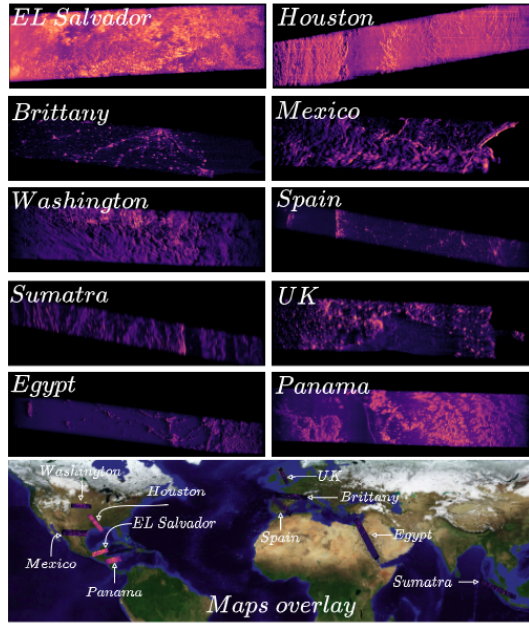


Figure 8. Motion-compensated maps. Top: The $H(\mathbf{u}', \boldsymbol{\theta})$ image using the motion parameters $\boldsymbol{\theta}$ estimated by our approach. Bottom: An overlay of motion-compensated maps over the earth map. The overlay was performed manually and the scale was a bit exaggerated to better see the motion-compensated maps.

noise heavily dominates the structure of the events. We evaluate our approach using data captured directly with an event camera on the ISS. The dataset comprises recordings that vary from 30 to 180 seconds and were captured under a diverse set of conditions. These include day/night recordings, different locations on earth, and varying weather conditions. The recordings contain an average of 7 million events per recording - See supplementary materials (Section I). The dataset does not come with an associated evaluation protocol and we, therefore, have defined an evaluation protocol based on RMS and a new metric called the Rate of Convergence (RoC). The RoC shows the rate of success of the optimisation algorithm in converging to the correct motion values. For simplicity, we used the Nelder-Mead optimisation (NMO) algorithm to search for the correct motion parameters and calculated the RoC by initialising the optimisation algorithm at every single point between $-30px/s$ and $30px/s$ and then calculating the overall percentage of how many times it successfully converges. The ground truth $\boldsymbol{\theta}$ was manually found for each recording.

3.1. Qualitative Results

Table 1 shows qualitative results using our proposed method compared with standard CMax [8]. Our approach always produces a single solution around the correct motion parameters, whereas CMax shows multiple extrema in each

case where the global maximum is much more prominent than the correct local maximum. A single correct motion solution indicates that the geometry of the IWE and the density of the events are no longer affecting the variance calculation. Our method also leads to better motion-compensated and sharp maps, which can be used for matching with existing satellite images and other orbital-related applications as in Figure 8.

3.2. Quantitative Results

In this section, we thoroughly examine the effectiveness of our proposed approach by comparing its performance with the existing CMax algorithm. We evaluate the results in terms of RMS and RoC. As shown in Table 2, our approach significantly outperforms CMax in every case.

The low RoC in CMax can be attributed to the presence of the global maximum, which often causes the NMO to be stuck at the global maximum and produce a high RMS error. Our method guarantees a single local maximum, which enables the NMO to converge successfully, resulting in a higher RoC.

Furthermore, the RoC of CMax is usually low since it only succeeds when the NMO is initialised close to the true motion parameters. In contrast, our approach is more robust and works effectively even when the global maximum is not dominant, as demonstrated in the Brittany case. Here, both methods converged since the recording had less noise compared to the rest of the data. This demonstrates the flexibility and general applicability of our approach to handling events with various density levels.

4. Conclusion

In this paper, we present an analytical solution to the noise-intolerance and the multiple extrema problems of the CMax framework. Our solution is purely based on geometrical principles and the physical properties of the events. First, we analysed these problems in 1D and 2D spaces by considering the events as a solid rectangle in 1D, and as a solid rectangular cuboid in 2D to demonstrate the influence of the change in geometries on the variance calculation. We then demonstrated how our analytical solution makes CMax invariant to the changes in the geometry and avoid having high contrast around the wrong motion parameters, without using any prior of the camera motion and regardless of the density of the events. The experimental results demonstrate the superior performance of our method compared to the state-of-the-art CMax when used on extremely noisy data.

Acknowledgement. This work was supported by AFOSR grant FA2386-23-1-4005 and USAFA grant FA7000-20-2-0009.




Table 1. Comparing the loss landscape produced by CMax [8] and our approach. Our approach shows a single solution corresponding to the correct motion θ , while CMax shows multiple maxima in each case. X-axis represents v_x [px/s] and Y-axis represents v_y [px/s] as a function of the variance (3), normalised between 0-1.

Table 2. Quantitative results. Final results of the ISS data comparing CMax [8] with our approach with RMS error and Rate of Convergence (RoC) as evaluation metrics.

	CMax [8]		Ours	
	<i>RMS</i>	<i>RoC%</i>	<i>RMS</i>	<i>RoC%</i>
EL Salvador	14.47	2.55	0.61	75.57
Houston	13.74	2.62	0.55	81.48
Brittany	0.08	83.13	0.01	83.57
Mexico	14.13	1.16	0.09	80.50
Washington	14.19	2.87	0.11	74.10
Spain	13.77	2.45	0.14	80.82
Sumatra	13.41	1.62	0.22	81.60
UK	12.84	2.02	0.28	82.89
Egypt	13.53	1.95	0.01	76.51
Panama	14.50	2.72	0.04	70.61

References

- [1] Saeed Afshar, Nicholas Ralph, Ying Xu, Jonathan Tapson, André van Schaik, and Gregory Cohen. Event-Based Feature Extraction Using Adaptive Selection Thresholds. *Sensors*, 20(6):1600, Mar. 2020. 2
- [2] Ignacio Alzugaray and Margarita Chli. Asynchronous Corner Detection and Tracking for Event Cameras in Real Time. *IEEE Robotics and Automation Letters*, 3(4):3177–3184, Oct. 2018. 2
- [3] Yeshwanth Bethi, Ying Xu, Gregory Cohen, Andre van Schaik, and Saeed Afshar. An optimised deep spiking neural network architecture without gradients. *IEEE Access*, 10:97912–97929, 2022. arXiv:2109.12813 [cs]. 2
- [4] Gregory Cohen, Saeed Afshar, Brittany Morreale, Travis Bessell, Andrew Wabnitz, Mark Rutten, and André van Schaik. Event-based Sensing for Space Situational Awareness. *The Journal of the Astronautical Sciences*, 66(2):125–141, June 2019. 2
- [5] Gregory Cohen, Saeed Afshar, and André van Schaik. Approaches for astrometry using event-based sensors. In *Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, Hawaii, USA*, 2018. 2
- [6] Thomas Finateu, Atsumi Niwa, Daniel Matolin, Koya Tsuchimoto, Andrea Mascheroni, Etienne Reynaud, Pooria Mostafalu, Frederick T. Brady, Ludovic Chotard, F. Legoff,

- Hirotsugu Takahashi, Hayato Wakabayashi, Yusuke Oike, and Christoph Posch. 5.10 a 1280×720 back-illuminated stacked temporal contrast event-based vision sensor with 4.86µm pixels, 1.066geps readout, programmable event-rate controller and compressive data-formatting pipeline. *2020 IEEE International Solid-State Circuits Conference - (ISSCC)*, pages 112–114, 2020. 1, 2
- [7] Guillermo Gallego, Mathias Gehrig, and Davide Scaramuzza. Focus is all you need: Loss functions for event-based vision. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12272–12281. 2, 3
- [8] Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3867–3876. IEEE. 1, 2, 3, 7, 8
- [9] Guillermo Gallego and Davide Scaramuzza. Accurate Angular Velocity Estimation With an Event Camera. *IEEE Robotics and Automation Letters*, 2(2):632–639, Apr. 2017. 2, 3
- [10] Mark Garcia. International space station facts and figures, 2023. Accessed on February 20, 2023. 2
- [11] Daniel Gehrig, Mathias Gehrig, Javier Hidalgo-Carrió, and Davide Scaramuzza. Video to events: Recycling video datasets for event cameras. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, June 2020. 2
- [12] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. EKLT: Asynchronous Photometric Feature Tracking Using Events and Frames. *International Journal of Computer Vision*, 128(3):601–618, Mar. 2020. 2
- [13] Suman Ghosh and Guillermo Gallego. Multi-Event-Camera Depth Estimation and Outlier Rejection by Refocused Events Fusion. *Advanced Intelligent Systems*, 4(12):2200221, Dec. 2022. 3
- [14] Arren Glover and Chiara Bartolozzi. Robust visual tracking with a freely-moving event camera. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3769–3776. IEEE. 2
- [15] Cheng Gu, Erik Learned-Miller, Daniel Sheldon, Guillermo Gallego, and Pia Bideau. The spatio-temporal poisson point process: A simple model for the alignment of event camera data. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [16] Jesse Hagenaars, Federico Paredes-Vallés, and Guido de Croon. Self-supervised learning of event-based optical flow with spiking neural networks. *Advances in Neural Information Processing Systems*, 34, 2021. 3
- [17] Craig Iaboni, Himanshu Patel, Deepan Lobo, Ji-Won Choi, and Pramod Abichandani. Event Camera Based Real-Time Detection and Tracking of Indoor Ground Robots. *IEEE Access*, 9:166588–166602, 2021. 2
- [18] Krzysztof Kamiński, Gregory Cohen, Tobi Delbruck, Michał Żołnowski, and Marcin Gdek. Observational evaluation of event cameras performance in optical space surveillance. 2019. 2
- [19] Hanme Kim, Ankur Handa, Ryad Benosman, Sio-Hoi Ieng, and Andrew Davison. Simultaneous mosaicing and tracking with an event camera. In *Proceedings of the British Machine Vision Conference 2014*, pages 26.1–26.12. British Machine Vision Association. 2, 3
- [20] Haram Kim and H. Jin Kim. Real-Time Rotational Motion Estimation With Contrast Maximization Over Globally Aligned Events. *IEEE Robotics and Automation Letters*, 6(3):6016–6023, July 2021. 3
- [21] Hanme Kim, Stefan Leutenegger, and Andrew J. Davison. Real-time 3d reconstruction and 6-dof tracking with an event camera. In *European Conference on Computer Vision*, 2016. 2
- [22] Beat Kueng, Elias Mueggler, Guillermo Gallego, and Davide Scaramuzza. Low-latency visual odometry using event-based feature tracks. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 16–23. IEEE. 2
- [23] 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 Transactions on Pattern Analysis and Machine Intelligence*, 39(7):1346–1359, July 2017. 2
- [24] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. A 128×128 120 dB 15 µs Latency Asynchronous Temporal Contrast Vision Sensor. *IEEE Journal of Solid-State Circuits*, 43(2):566–576, 2008. 1
- [25] Daqi Liu, Alvaro Parra, and Tat-Jun Chin. Globally optimal contrast maximisation for event-based motion estimation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6348–6357. IEEE. 2, 3
- [26] Daqi Liu, Alvaro Parra, and Tat-Jun Chin. Spatiotemporal registration for event-based visual odometry. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4935–4944. IEEE. 2
- [27] Florian Mähknecht, Daniel Gehrig, Jeremy Nash, Friedrich M. Rockenbauer, Benjamin Morrell, Jeff Delaune, and Davide Scaramuzza. Exploring Event Camera-Based Odometry for Planetary Robots. *IEEE Robotics and Automation Letters*, 7(4):8651–8658, Oct. 2022. 2
- [28] Ana I. Maqueda, Antonio Loquercio, Guillermo Gallego, Narciso Garcia, and Davide Scaramuzza. Event-based vision meets deep learning on steering prediction for self-driving cars. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5419–5427. IEEE. 2
- [29] Matthew G. McHarg, Richard L. Balthazor, Brian J. McReynolds, David H. Howe, Colin J. Maloney, Daniel O’Keefe, Rayomand Bam, Gabriel Wilson, Paras Karki, Alexandre Marcireau, and Gregory Cohen. Falcon Neuro: an event-based sensor on the International Space Station. *Optical Engineering*, 61(08), Aug. 2022. 2
- [30] Sofia McLeod, Gabriele Meoni, Dario Izzo, Anne Mergy, Daqi Liu, Yasir Latif, Ian Reid, and Tat-Jun Chin. Globally Optimal Event-Based Divergence Estimation for Ventral Landing. In Leonid Karlinsky, Tomer Michaeli, and Ko Nishino, editors, *Computer Vision – ECCV 2022 Workshops*, volume 13801, pages 3–20. Springer Nature Switzerland, Cham, 2023. Series Title: Lecture Notes in Computer Science. 2

- [31] Moritz B. Milde, Hermann Blum, Alexander Dietmüller, Dora Sumislawska, Jörg Conradt, Giacomo Indiveri, and Yulia Sandamirskaya. Obstacle Avoidance and Target Acquisition for Robot Navigation Using a Mixed Signal Analog/Digital Neuromorphic Processing System. *Frontiers in Neurobotics*, 11:28, July 2017. 2
- [32] Elias Mueggler. *Event-based vision for high-speed robotics*. PhD thesis, University of Zurich, 2017. 2
- [33] Elias Mueggler, Basil Huber, and Davide Scaramuzza. Event-based, 6-DOF pose tracking for high-speed maneuvers. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2761–2768. IEEE, 2
- [34] Zhenjiang Ni, Sio-Hoi Ieng, Christoph Posch, Stéphane Régnier, and Ryad Benosman. Visual Tracking Using Neuromorphic Asynchronous Event-Based Cameras. *Neural Computation*, 27(4):925–953, Apr. 2015. 2
- [35] Urbano Miguel Nunes and Yiannis Demiris. Robust Event-Based Vision Model Estimation by Dispersion Minimisation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9561–9573, Dec. 2022. 2, 3
- [36] Garrick Orchard, Cedric Meyer, Ralph Etienne-Cummings, Christoph Posch, Nitish Thakor, and Ryad Benosman. HFirst: A Temporal Approach to Object Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(10):2028–2040, Oct. 2015. 2
- [37] Takehiro Ozawa, Yusuke Sekikawa, and Hideo Saito. Accuracy and Speed Improvement of Event Camera Motion Estimation Using a Bird’s-Eye View Transformation. *Sensors*, 22(3):773, Jan. 2022. 2
- [38] Takehiro Ozawa, Yusuke Sekikawa, and Hideo Saito. Recursive Contrast Maximization for Event-Based High-Frequency Motion Estimation. *IEEE Access*, 10:125376–125386, 2022. 2
- [39] Chethan M. Parameshwara, Nitin J. Sanket, Chahat Deep Singh, Cornelia Fermüller, and Yiannis Aloimonos. 0-MMS: Zero-shot multi-motion segmentation with a monocular event camera. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9594–9600. IEEE, 3
- [40] Xin Peng, Ling Gao, Yifu Wang, and Laurent Kneip. Globally-Optimal Contrast Maximisation for Event Cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2021. 2, 3
- [41] Nicholas Ralph, Damien Joubert, Andrew Jolley, Saeed Afshar, Nicholas Tothill, André van Schaik, and Gregory Cohen. Real-Time Event-Based Unsupervised Feature Consolidation and Tracking for Space Situational Awareness. *Frontiers in Neuroscience*, 16:821157, May 2022. 2
- [42] Nicholas Owen Ralph, Alexandre Marcireau, Saeed Afshar, Nicholas Tothill, André van Schaik, and Gregory Cohen. Astrometric Calibration and Source Characterisation of the Latest Generation Neuromorphic Event-based Cameras for Space Imaging. Technical Report arXiv:2211.09939, arXiv, Nov. 2022. arXiv:2211.09939 [astro-ph] type: article. 2
- [43] Bharath Ramesh, S. Zhang, Zhi Wei Lee, Zhi Gao, G. Orchard, and Cheng Xiang. Long-term object tracking with a moving event camera. In *British Machine Vision Conference*, 2018. 2
- [44] Henri Rebecq, Guillermo Gallego, Elias Mueggler, and Davide Scaramuzza. EMVS: Event-based multi-view stereo—3D reconstruction with an event camera in real-time. *Int. J. Comput. Vis.*, 126:1394–1414, Dec. 2018. 2
- [45] Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. EMVS: Event-based Multi-View Stereo. In *Proceedings of the British Machine Vision Conference 2016*, pages 63.1–63.11, York, UK, 2016. British Machine Vision Association. 3
- [46] Christian Reinbacher, Gottfried Munda, and Thomas Pock. Real-time panoramic tracking for event cameras. In *2017 IEEE International Conference on Computational Photography (ICCP)*, pages 1–9. IEEE, 3
- [47] Hochang Seok and Jongwoo Lim. Robust feature tracking in DVS event stream using bézier mapping. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1647–1656. IEEE, 2, 3
- [48] Shintaro Shiba, Yoshimitsu Aoki, and Guillermo Gallego. Event Collapse in Contrast Maximization Frameworks. *Sensors*, 22(14):5190, July 2022. arXiv:2207.04007 [cs, math]. 2
- [49] Shintaro Shiba, Yoshimitsu Aoki, and Guillermo Gallego. Secrets of event-based optical flow. In *European Conference on Computer Vision (ECCV)*, pages 628–645, 2022. 3
- [50] Shintaro Shiba, Yoshimitsu Aoki, and Guillermo Gallego. A Fast Geometric Regularizer to Mitigate Event Collapse in the Contrast Maximization Framework. *Advanced Intelligent Systems*, 5(3):2200251, Mar. 2023. 2, 3
- [51] Olaf Sikorski, Dario Izzo, and Gabriele Meoni. Event-based spacecraft landing using time-to-contact. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1941–1950. IEEE, 2
- [52] Amos Sironi, Manuele Brambilla, Nicolas Bourdis, Xavier Lagorce, and Ryad Benosman. HATS: Histograms of averaged time surfaces for robust event-based object classification. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1731–1740. IEEE, 2
- [53] Timo Stoffregen, Guillermo Gallego, Tom Drummond, Lindsay Kleeman, and Davide Scaramuzza. Event-based motion segmentation by motion compensation. In *Int. Conf. Comput. Vis. (ICCV)*, 2019. 3
- [54] Timo Stoffregen and Lindsay Kleeman. Event cameras, contrast maximization and reward functions: An analysis. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12292–12300. IEEE, 2, 3
- [55] David Tedaldi, Guillermo Gallego, Elias Mueggler, and Davide Scaramuzza. Feature detection and tracking with the dynamic and active-pixel vision sensor (DAVIS). In *2016 Second International Conference on Event-based Control, Communication, and Signal Processing (EBCCSP)*, pages 1–7. IEEE, 2
- [56] Antoni Rosinol Vidal, Henri Rebecq, Timo Horstschaefer, and Davide Scaramuzza. Ultimate SLAM? Combining Events, Images, and IMU for Robust Visual SLAM in HDR and High-Speed Scenarios. *IEEE Robotics and Automation Letters*, 3(2):994–1001, Apr. 2018. 2
- [57] Antonio Vitale, Alpha Renner, Celine Nauer, Davide Scaramuzza, and Yulia Sandamirskaya. Event-driven vision and

- control for uavs on a neuromorphic chip. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 103–109. IEEE, 2021. 2
- [58] David Weikersdorfer, Raoul Hoffmann, and Jörg Conradt. Simultaneous localization and mapping for event-based vision systems. In Mei Chen, Bastian Leibe, and Bernd Neumann, editors, *Computer Vision Systems*, volume 7963, pages 133–142. Springer Berlin Heidelberg. Series Title: Lecture Notes in Computer Science. 2
- [59] Christian E. Willert and Joachim Klinner. Event-based imaging velocimetry: an assessment of event-based cameras for the measurement of fluid flows. *Experiments in Fluids*, 63(6):101, June 2022. 2
- [60] Feihu Zhang, Yaohui Zhong, Liyuan Chen, and Zhiliang Wang. Event-Based Circular Detection for AUV Docking Based on Spiking Neural Network. *Frontiers in Neuro-robotics*, 15:815144, Jan. 2022. 2
- [61] Zelin Zhang, Anthony Yezzi, and Guillermo Gallego. Formulating Event-based Image Reconstruction as a Linear Inverse Problem with Deep Regularization using Optical Flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–18, 2022. 3
- [62] Yi Zhou, Guillermo Gallego, Xiuyuan Lu, Siqi Liu, and Shaojie Shen. Event-based motion segmentation with spatio-temporal graph cuts. *IEEE Transactions on Neural Network and Learning Systems*, 2021. 3
- [63] Alex Zihao Zhu, Nikolay Atanasov, and Kostas Daniilidis. Event-based feature tracking with probabilistic data association. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4465–4470. IEEE. 2
- [64] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Unsupervised event-based learning of optical flow, depth, and egomotion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 989–997, 2019. 3