



Globally-Optimal Event Camera Motion Estimation

Xin Peng^{1,2}, Yifu Wang³, Ling Gao¹, and Laurent Kneip^{1,4}(✉)

¹ Mobile Perception Lab, SIST, ShanghaiTech University, Shanghai, China
kneip.laurent@gmail.com

² Shanghai Institute of Microsystems and Information Technology,
Chinese Academy of Sciences, Beijing, China

³ Australian National University, Canberra, Australia

⁴ Shanghai Engineering Research Center of Intelligent Vision and Imaging,
Shanghai, China

Abstract. Event cameras are bio-inspired sensors that perform well in HDR conditions and have high temporal resolution. However, different from traditional frame-based cameras, event cameras measure asynchronous pixel-level brightness changes and return them in a highly discretised format, hence new algorithms are needed. The present paper looks at fronto-parallel motion estimation of an event camera. The flow of the events is modeled by a general homographic warping in a space-time volume, and the objective is formulated as a maximisation of contrast within the image of unwarped events. However, in stark contrast to prior art, we derive a globally optimal solution to this generally non-convex problem, and thus remove the dependency on a good initial guess. Our algorithm relies on branch-and-bound optimisation for which we derive novel, recursive upper and lower bounds for six different contrast estimation functions. The practical validity of our approach is supported by a highly successful application to AGV motion estimation with a downward facing event camera, a challenging scenario in which the sensor experiences fronto-parallel motion in front of noisy, fast moving textures.

Keywords: Event cameras · Motion estimation · Contrast maximisation · Global optimality · Branch and bound

1 Introduction

Camera motion estimation is an important technology with many applications in automation, smart transportation, and assistive technologies. However, despite the fact that a certain level of maturity has already been reached, we keep facing

X. Peng, Y. Wang and L. Gao—indicates equal contribution.

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-58574-7_4) contains supplementary material, which is available to authorized users.

challenges in scenarios with high dynamics, low texture distinctiveness, or challenging illumination conditions [5, 9]. Event cameras—also called dynamic vision sensors—present an interesting alternative in this regard, as they pair HDR with high temporal resolution. The potential advantages and challenges behind event-based vision are well explained by the original work of Brandli et al. [3] as well as the recent survey by Gallego et al. [10].

Our work considers fronto-parallel motion estimation of an event camera. The flow of the events is hereby modelled by a general homographic warping in a space-time volume, and motion may be estimated by maximisation of contrast in the image of unwarped events [12]. Various reward functions that maximise contrast have been presented and analysed in the recent works of Gallego et al. [11] and Stoffregen and Kleeman [29], and successfully used for solving a variety of problems with event cameras such as optical flow [12, 28, 32, 34, 36, 37], segmentation [21, 27, 28], 3D reconstruction [26, 32, 35, 37], and motion estimation [12, 13]. Our work focuses on the latter problem of camera motion estimation. However—different from many of the aforementioned works—we propose the first globally optimal solution to the underlying contrast maximisation problem, an important point given its generally non-convex nature.

Our detailed contributions are as follows:

- We solve the global maximisation of contrast functions via Branch and Bound.
- We derive bounds for six different contrast estimation functions. The bounds are furthermore calculated recursively, which enables efficient processing.
- We successfully apply this strategy to Autonomous Ground Vehicle (AGV) planar motion estimation with a downward facing event camera (cf. Fig. 1), a problem that is complicated by motion blur, challenging illumination conditions, and indistinctive, noisy textures. We prove that using an event camera can solve these challenges, hence outperforming alternatives given by regular cameras.




Fig. 1. (a): AGV equipped with a downward facing event camera for vehicle motion estimation. (b)–(d): collected image with detectable corners, image of warped events with $\theta = 0$, and image of warped events with optimal parameters $\hat{\theta}$.

2 Contrast Maximisation

Gallego et al. [12] recently introduced contrast maximisation as a unifying framework allowing the solution of several important problems for dynamic vision sensors, in particular motion estimation problems in which the effect of camera motion may be described by a homography (e.g. motion in front of a plane, pure rotation). Our work relies on contrast maximisation, which we therefore briefly review in the following.

An event camera outputs a sequence of *events* denoting temporal logarithmic brightness changes above a certain threshold. An event $e = \{\mathbf{x}, t, s\}$ is described by its pixel position $\mathbf{x} = [x \ y]^T$, timestamp t , and polarity s (the latter indicates whether the brightness is increasing or decreasing, and is ignored in the present work). The core idea of contrast maximisation is relatively straightforward: The flow of the events is modelled by a time-parametrised homography. Given its position and time-stamp, every event may therefore be warped back along a point-trajectory into a reference view with timestamp t_{ref} . Since events are more likely to be generated by high-gradient edges, correct homographic warping parameters will likely lead to a sharp Image of Warped Events (IWE) in which events align along a crisp edge-map. Gallego et al. [12] simply propose to consider the contrast of the IWE as a reward function to identify the correct homographic warping parameters. Note that homographic warping functions include 2D affine and Euclidean transformations, and thus can be used in a variety of vision problems such as optical flow, feature tracking, or fronto-parallel motion estimation.

Suppose we are given a set of N events $\mathcal{E} = \{e_k\}_{k=1}^N$. We define a general warping function $\mathbf{x}'_k = W(\mathbf{x}_k, t_k; \boldsymbol{\theta})$ that returns the position \mathbf{x}'_k of an event e_k in the reference view at time t_{ref} . $\boldsymbol{\theta}$ is a vector of warping parameters. The IWE is generated by accumulating warped events at each discrete pixel location:

$$I(\mathbf{p}_{ij}; \boldsymbol{\theta}) = \sum_{k=1}^N \mathbf{1}(\mathbf{p}_{ij} - \mathbf{x}'_k) = \sum_{k=1}^N \mathbf{1}(\mathbf{p}_{ij} - W(\mathbf{x}_k, t_k; \boldsymbol{\theta})), \quad (1)$$

where $\mathbf{1}(\cdot)$ is an indicator function that counts 1 if the absolute value of $(\mathbf{p}_{ij} - \mathbf{x}'_k)$ is less than a threshold ϵ in each coordinate, and otherwise 0. \mathbf{p}_{ij} is a pixel in the IWE with coordinates $[i \ j]^T$, and we refer to it as an *accumulator* location. We set $\epsilon = 0.5$ such that each warped event will increment one accumulator only.

Existing approaches replace the indicator function with a Gaussian kernel to make the IWE a smooth function of the warped events, and thus solve contrast maximisation problems via local optimisation methods (cf. [11–13]). In contrast, we show how our proposed method is able to find the global optimum of the above, discrete objective function.

As introduced in [11, 29], reward functions for event un-warping all rely on the idea of maximising the contrast or sharpness of the IWE (they have also been denoted as *focus loss functions*). They proceed by integration over the entire set of accumulators, which we denote \mathcal{P} . The most relevant ones for us are summarized in Table 1. Note that for L_{Var} , μ_I is the mean value of $I(\mathbf{p}_{ij}; \boldsymbol{\theta})$ over

all pixels (a function of θ itself), and N_p is the total number of accumulators in I . For L_{SoSA} , δ is a design parameter called the *shift factor*. Different from other objectives functions, locations with few accumulations will contribute more to L_{SoSA} . The intuition here is that more empty locations again mean more events that are concentrated at fewer accumulators. L_{SoEaS} is a combination of L_{SoS} and L_{SoE} . Similarly, L_{SoSAaS} is a combination of L_{SoS} and L_{SoSA} .

Let us now proceed to the main contribution of our work, which is a derivation of bounds on the above objectives as required by Branch and Bound.

Table 1. Contrast functions evaluated in this work

Sum of Squares (SoS)	$L_{\text{SoS}}(\theta) = \sum_{\mathbf{p}_{ij} \in \mathcal{P}} I(\mathbf{p}_{ij}; \theta)^2$
Variance (Var)	$L_{\text{Var}}(\theta) = \frac{1}{N_p} \sum_{\mathbf{p}_{ij} \in \mathcal{P}} (I(\mathbf{p}_{ij}; \theta) - \mu_I)^2$
Sum of Exponentials (SoE)	$L_{\text{SoE}}(\theta) = \sum_{\mathbf{p}_{ij} \in \mathcal{P}} e^{I(\mathbf{p}_{ij}; \theta)}$
Sum of Suppressed Accumulations (SoSA)	$L_{\text{SoSA}}(\theta) = \sum_{\mathbf{p}_{ij} \in \mathcal{P}} e^{-I(\mathbf{p}_{ij}; \theta) \cdot \delta}$
SoE and Squares (SoEaS)	$L_{\text{SoEaS}}(\theta) = \sum_{\mathbf{p}_{ij} \in \mathcal{P}} I(\mathbf{p}_{ij}; \theta)^2 + e^{I(\mathbf{p}_{ij}; \theta)}$
SoSA and Squares (SoSAaS)	$L_{\text{SoSAaS}}(\theta) = \sum_{\mathbf{p}_{ij} \in \mathcal{P}} I(\mathbf{p}_{ij}; \theta)^2 + e^{-I(\mathbf{p}_{ij}; \theta) \cdot \delta}$

3 Globally Maximised Contrast Using Branch and Bound

Figure 2 illustrates how contrast maximisation for motion estimation is in general a non-convex problem, meaning that local optimisation may be sensitive to the initial parameters and not find the global optimum. We tackle this problem by introducing a globally optimal solution to contrast maximisation using Branch and Bound (BnB) optimisation. BnB is an algorithmic paradigm in which the solution space is subdivided into branches in which we then find upper and lower bounds for the maximal objective value. The globally optimal solution is isolated by an iterative search in which entire branches are discarded if their upper bound for the maximum objective value remains lower than the corresponding lower bound in another branch. The most important factor deciding the effectiveness of this approach is given by the tightness of the bounds.

Our core contribution is given by a recursive method to efficiently calculate upper and lower bounds for the maximum value of a contrast maximisation function over a given branch. In short, the main idea is given by expressing a bound over $(N + 1)$ events as a function of the bound over N events plus the contribution of one additional event. The strategy can be similarly applied to all six aforementioned contrast functions, which is why we limit the exposition to the derivation of bounds for L_{SoS} . Detailed derivations for all loss functions are provided in the supplementary material.




Fig. 2. Visualization of the Sum of Squares contrast function. The camera is moving in front of a plane, and the motion parameters are given by translational and rotational velocity (cf. Sect. 4). The sub-figures from left to right are functions with increasing Noise-to-Events (N/E) ratios. Note that contrast functions are non-convex.

3.1 Objective Function

In the following, we assume that $L = L_{\text{SoS}}$. The maximum objective function value over all N events in a given time interval $[t_{\text{ref}}, t_{\text{ref}} + \Delta T]$ is given by

$$L_N = \max_{\theta \in \Theta} \sum_{\mathbf{p}_{ij} \in \mathcal{P}} \left[\sum_{k=1}^N \mathbf{1}(\mathbf{p}_{ij} - W(\mathbf{x}_k, t_k; \theta)) \right]^2, \quad (2)$$

where Θ is the search space (i.e. branch or sub-branch) over which we want to maximise the objective. Most globally optimal methods for geometric computer vision problems find bounds by a spatial division of the problem into individual, simpler maximisation sub-problems (cf. [6]). However, the contrast maximisation objective is related to the distribution over the entire IWE and not just individual accumulators, which complicates this strategy.

3.2 Upper and Lower Bound

The bounds are calculated recursively by processing the events and one-by-one, each time updating the IWE. The event are notably processed in temporal order with increasing timestamps.

For the lower bound, it is readily given by evaluating the contrast function at an arbitrary point on the interval Θ , which is commonly picked as the interval center θ_0 . We present a recursive rule to efficiently evaluate the lower bound.

Theorem 1. *For search space Θ centered at θ_0 , the lower bound of SoS-based contrast maximisation may be given by*

$$\underline{L}_{N+1} = \underline{L}_N + 1 + 2I^N(\eta_{N+1}^{\theta_0}; \theta_0), \quad (3)$$

where $I^N(\mathbf{p}_{ij}; \theta_0)$ is the incrementally constructed IWE, its exponent N denotes the number of events that have already been taken into account, and

$$\eta_{N+1}^{\theta_0} = \text{round}(W(\mathbf{x}_{N+1}, t_{N+1}; \theta_0)) \quad (4)$$

returns the accumulator closest to the warped position of the $(N + 1)$ -th event.

Proof. According to the definition of sum of the square focus loss function,

$$\begin{aligned}
\underline{L}_{N+1} &= \sum_{\mathbf{p}_{ij} \in \mathcal{P}} \left[\sum_{k=1}^{N+1} \mathbf{1}(\mathbf{p}_{ij} - W(\mathbf{x}_k, t_k; \boldsymbol{\theta}_0)) \right]^2 \\
&= \sum_{\mathbf{p}_{ij} \in \mathcal{P}} [I^N(\mathbf{p}_{ij}; \boldsymbol{\theta}_0) + \mathbf{1}(\mathbf{p}_{ij} - W(\mathbf{x}_{N+1}, t_{N+1}; \boldsymbol{\theta}_0))]^2 \quad (5) \\
&= a + b + c, \text{ where} \\
a &= \sum_{\mathbf{p}_{ij} \in \mathcal{P}} I^N(\mathbf{p}_{ij}; \boldsymbol{\theta}_0)^2, \\
b &= 2 \sum_{\mathbf{p}_{ij} \in \mathcal{P}} [\mathbf{1}(\mathbf{p}_{ij} - W(\mathbf{x}_{N+1}, t_{N+1}; \boldsymbol{\theta}_0)) I^N(\mathbf{p}_{ij}; \boldsymbol{\theta}_0)], \\
c &= \sum_{\mathbf{p}_{ij} \in \mathcal{P}} [\mathbf{1}(\mathbf{p}_{ij} - W(\mathbf{x}_{N+1}, t_{N+1}; \boldsymbol{\theta}_0))]^2.
\end{aligned}$$

It is clear that $a = \underline{L}_N$. In c , owing to the definition of our indicator function, only the \mathbf{p}_{ij} which is closest to $W(\mathbf{x}_{N+1}, t_{N+1}; \boldsymbol{\theta}_0)$ makes a contribution, thus we have $c = 1$. For b , the term $\mathbf{1}(\mathbf{p}_{ij} - W(\mathbf{x}_{N+1}, t_{N+1}; \boldsymbol{\theta}_0))$ is simply zero unless we are considering an accumulator $\mathbf{p}_{ij} = \boldsymbol{\eta}_{N+1}^{\theta_0}$, which gives $b = 2I^N(\boldsymbol{\eta}_{N+1}^{\theta_0}; \boldsymbol{\theta}_0)$. Thus we obtain (3). Note that the IWE is iteratively updated by incrementing the accumulator which locates closest to $\boldsymbol{\eta}_{N+1}^{\theta_0}$.




Fig. 3. (a) Incremental update of the IWE. For each new event e , we choose and increment the currently maximal accumulator in the bounding box \mathcal{P}^θ around all possible locations $W(\mathbf{x}, t; \boldsymbol{\theta} \in \Theta)$. We simply increment the center of the bounding box if no other accumulator exists. (b) Bounding boxes of two temporally distinct events generated by the same point in 3D.

We now proceed to our main contribution, a recursive upper bound for the contrast maximisation problem. Let us define \mathcal{P}_i^θ as the bounding box around all possible locations $W(\mathbf{x}_i, t_i; \boldsymbol{\theta} \in \Theta)$ of the un-warped event. Lemma 1 is introduced as follows.

Lemma 1. *Given a search space $\boldsymbol{\theta} \in \Theta$, for a small enough time interval, if $W(\mathbf{x}_i, t_i; \boldsymbol{\theta}) = W(\mathbf{x}_j, t_j; \boldsymbol{\theta})$ and $0 < i < j \leq N$, we have $\mathcal{P}_i^\theta \subseteq \mathcal{P}_j^\theta$. An intuitive explanation is given in Fig. 3(b).*

Lemma 1 now enables us to derive our recursive upper bound.

Theorem 2. *The upper bound of the objective function L_N for SoS-based contrast maximisation satisfies*

$$L_{N+1} = L_N + 1 + 2I^N(\boldsymbol{\eta}_{N+1}^{\hat{\theta}}; \hat{\theta}) \quad (6)$$

$$\leq \overline{L}_N + 1 + 2Q^N = \overline{L}_{N+1}, \quad (7)$$

$$\text{where } Q^N = \max_{\mathbf{p}_{ij} \in \mathcal{P}_{N+1}^{\Theta}} \overline{I}^N(\mathbf{p}_{ij}) \geq I^N(\boldsymbol{\eta}_{N+1}^{\hat{\theta}}; \hat{\theta})$$

$\mathcal{P}_{N+1}^{\Theta}$ is a bounding box for the $(N+1)$ -th event. $\hat{\theta}$ is the optimal parameter set that maximises L_{N+1} over the interval Θ . $\overline{I}^N(\mathbf{p}_{ij})$ is the value of pixel \mathbf{p}_{ij} in the upper bound IWE, a recursively constructed image in which we always increment the maximum accumulator within the bounding box $\mathcal{P}_{N+1}^{\Theta}$ (i.e. the one that we used to define the value of Q^N). The incremental construction of $\overline{I}^N(\mathbf{p}_{ij})$ is illustrated in Fig. 3(a).

Proof. (6) is straightforwardly derived from (3). The proof of inequation (7) then proceeds by mathematical induction.

For $N = 0$, it is obvious that $L_0 = \overline{L}_0 = 0$. Similarly, for $N = 1$, $L_1 = 1 \leq \overline{L}_0 + 1 + 0$, and $Q^0 = I^0(\boldsymbol{\eta}_1^{\hat{\theta}}; \hat{\theta}) = 0$ (which satisfies Theorem 2). We now assume that \overline{L}_n as well as the corresponding upper bound IWE \overline{I}^n are given for all $0 < n \leq N$. We furthermore assume that they satisfy Theorem 2. Our aim is to prove that (7) holds for the $(N+1)$ -th event. It is clear that $\overline{L}_N \geq L_N$, and we only need to prove that $Q^N \geq I^N(\boldsymbol{\eta}_{N+1}^{\hat{\theta}}; \hat{\theta})$, for which we will make use of Lemma 1. There are two cases to be distinguished:

- The first case is if there exists an event ϵ_k with $0 < k < N+1$ and for which $\boldsymbol{\eta}_k^{\hat{\theta}} = \boldsymbol{\eta}_{N+1}^{\hat{\theta}}$. In other words, the k -th and the $(N+1)$ -th event are warped to a same accumulator if choosing the locally optimal parameters. Note that if there are multiple previous events for which this condition holds, the k -th event is chosen to be the most recent one. Given our assumptions, \overline{L}_{k-1} as well as the $(k-1)$ -th constructed upper bound IWE satisfy Theorem 2, which means that $Q^{k-1} \geq I^{k-1}(\boldsymbol{\eta}_k^{\hat{\theta}}; \hat{\theta})$. Let $\mathbf{p}_k \in \mathcal{P}_k^{\Theta}$ now be the pixel location with maximum intensity in $\overline{I}^{k-1}(\mathbf{p}_k)$. Then, the k -th updated IWE satisfies $\overline{I}^k(\mathbf{p}_k) = Q^{k-1} + 1 \geq I^{k-1}(\boldsymbol{\eta}_k^{\hat{\theta}}; \hat{\theta}) + 1$. According to Lemma 1, we have $\mathcal{P}_k^{\Theta} \subseteq \mathcal{P}_{N+1}^{\Theta}$, therefore $\mathbf{p}_k \subseteq \mathcal{P}_{N+1}^{\Theta}$, and $Q^N \geq \overline{I}^k(\mathbf{p}_k) \geq I^{k-1}(\boldsymbol{\eta}_k^{\hat{\theta}}; \hat{\theta}) + 1$. With optimal warp parameters $\hat{\theta}$, events with indices from $k+1$ to N will not locate at $\boldsymbol{\eta}_{N+1}^{\hat{\theta}}$, and therefore $I^{k-1}(\boldsymbol{\eta}_k^{\hat{\theta}}; \hat{\theta}) + 1 = I^N(\boldsymbol{\eta}_{N+1}^{\hat{\theta}}; \hat{\theta}) \leq Q^N$.
- If there is no such a event, it is obvious that $Q^N \geq I^N(\boldsymbol{\eta}_{N+1}^{\hat{\theta}}; \hat{\theta})$.

With the basic cases and the induction step proven, we conclude our proof that Theorem 2 holds for all natural numbers N .

We apply the proposed strategy to derive upper and lower bounds for all six aforementioned contrast functions, and list them in Table 2. Note that the initial case varies for different loss functions. The globally-optimal contrast maximisation framework (GOCMF) is outlined in Algorithm 1 and Algorithm 2. We propose a nested strategy for calculating upper bounds, in which the outer layer *RB* evaluates the objective function, while the inner layer *BB* estimates the bounding box \mathcal{P}_N^Θ and depends on the specific motion parametrisation.

Table 2. Recursive Upper and Lower Bounds

	Upper bound \overline{L}_N	Lower bound \underline{L}_N	L_0
SoS	$\overline{L}_{N-1} + 1 + 2Q$	$\underline{L}_{N-1} + 1 + 2I^{N-1}(\eta_N^{\theta_0}; \theta_0)$	0
Var	$\overline{L}_{N-1} + \frac{1}{N_p} - \frac{2\mu_I}{N_p} + \frac{2}{N_p}Q$	$\underline{L}_{N-1} + \frac{1}{N_p} - \frac{2\mu_I}{N_p} + \frac{2}{N_p}I^{N-1}(\eta_N^{\theta_0}; \theta_0)$	μ_I^2
SoE	$\overline{L}_{N-1} + (e-1)e^Q$	$\underline{L}_{N-1} + (e-1)e^{I^{N-1}(\eta_N^{\theta_0}; \theta_0)}$	N_p
SoSA	$\overline{L}_{N-1} + (e^{-\delta} - 1)e^{-\delta \cdot Q}$	$\underline{L}_{N-1} + (e^{-\delta} - 1)e^{-\delta \cdot I^{N-1}(\eta_N^{\theta_0}; \theta_0)}$	N_p
SoEaS	$\overline{L}_{N-1} + 1 + 2Q + (e-1)e^Q$	$\underline{L}_{N-1} + 1 + 2I^{N-1}(\eta_N^{\theta_0}; \theta_0) + (e-1)e^{I^{N-1}(\eta_N^{\theta_0}; \theta_0)}$	N_p
SoSAaS	$\overline{L}_{N-1} + 1 + 2Q + (e^{-\delta} - 1)e^{-\delta Q}$	$\underline{L}_{N-1} + 1 + 2I^{N-1}(\eta_N^{\theta_0}; \theta_0) + (e^{-\delta} - 1)e^{-\delta I^{N-1}(\eta_N^{\theta_0}; \theta_0)}$	N_p

Algorithm 1. GOCMF: globally optimal contrast maximisation framework

Input: event set \mathcal{E} , initial search space Θ , branching limit N_b

Output: optimal warping parameters $\hat{\theta}$

- 1: Initialize $\hat{\theta}$ with the center of Θ ,
- 2: $L^* \leftarrow 0$, $S \leftarrow \{\text{RB}(\mathcal{E}, \Theta), \Theta\}$
- 3: Push S into queue Q , $S^* \leftarrow S$
- 4: **while** $i < N_b$ **do**
- 5: $L^* \leftarrow 0$
- 6: **if** $S^*.\underline{L} == S^*.\overline{L}$ **then**
- 7: $\hat{\theta} \leftarrow$ Center of $S^*.\Theta$, **break**
- 8: **for** each node $S \in Q$ **do**
- 9: Pop S , split into subspaces S_j
- 10: **for** all subspaces S_j **do**
- 11: $\{S_j.\underline{L}, S_j.\overline{L}\} \leftarrow \text{RB}(\mathcal{E}, \Theta_j)$
- 12: **if** $S_j.\underline{L} > L^*$ **then**
- 13: $L^* \leftarrow S_j.\underline{L}$, $S^* \leftarrow S_j$
- 14: Push S_j into Q
- 15: Prune branches in Q
- 16: $i \leftarrow i + 1$
- 17: **return** $\hat{\theta}$

Algorithm 2. RB: recursive bounds calculation

Input: event set \mathcal{E} , search space Θ

Output: lower bound \underline{L} , upper bound \overline{L}

- 1: Initialize accumulator images \overline{I} and I with zeros
- 2: Initialize \underline{L} , \overline{L} according to Table 2
- 3: $\theta_0 \leftarrow$ center of Θ
- 4: **for** each event $e_k \in \mathcal{E}$ **do**
- 5: $\mathcal{P}_k^\Theta \leftarrow \text{BB}(W(\cdot), \Theta, e_k)$
- 6: $Q \leftarrow \max_{\mathbf{p}_{ij} \in \mathcal{P}_k^\Theta} \overline{I}(\mathbf{p}_{ij})$
- 7: $\eta_k^{\theta_0} \leftarrow \text{round}(W(\mathbf{x}_k, t_k; \theta_0))$
- 8: Update \underline{L} , \overline{L} (cf. Table 2)
- 9: $\nu_k \leftarrow \text{argmax}_{\mathbf{p}_{ij} \in \mathcal{P}_k^\Theta} \overline{I}(\mathbf{p}_{ij})$
- 10: $\overline{I}(\nu_k) \leftarrow \overline{I}(\nu_k) + 1$
- 11: $I(\eta_k^{\theta_0}) \leftarrow I(\eta_k^{\theta_0}) + 1$
- 12: **return** \underline{L} , \overline{L}

4 Application to Visual Odometry with a Downward-Facing Event Camera

Motion estimation for planar Autonomous Ground Vehicles (AGVs) is an important problem in intelligent transportation [17, 24, 30]. An interesting alternative is given by employing a downward instead of a forward facing camera, thus permitting direct observation of the ground plane with known depth. This largely simplifies the geometry of the problem and notably turns the image-to-image warping into a homographic mapping that is linear in homogeneous space. The strategy is widely used in relevant applications such as sweeping robots and factory AGVs, and a good review is presented in [1]. However, the method is affected by potentially severe challenges given by the image appearance: a) reliable feature matching or even extraction may be difficult for certain noisy ground textures, b) fast motion may easily lead to motion blur, and c) stable appearance may require artificial illumination. Many existing methods therefore do not employ feature correspondences but aim at a correspondence-less alignment or even a full photometric image alignment. Besides more classical RANSAC-based hypothesise-and-test schemes [7], the community therefore has also developed appearance-based template matching approaches [8, 15, 22, 23, 33], solvers based on efficient second-order minimisation [18, 20, 38], and methods exploiting the Fast Fourier Transform [2, 25], the Fourier-Mellin Transform [16, 19], or the Improved Fourier Mellin Invariant [4, 31]. In an attempt to tackle highly self-similar ground textures, Dille et al. [8] propose to use an optical flow sensor instead of a regular CMOS camera.

A critical question is given by the position of the camera. The camera may hang in the front or rear of the vehicle, which gives increased distance to the ground plane and in turn reduces motion blur. However, it also causes moving shadows in the image, and generally complicates the stabilisation of the image appearance and thus repeatable feature detection or region-based matching. A common alternative therefore is given by installing the camera underneath the vehicle paired with an artificial light source (e.g. [2, 8]). However, the short distance to the ground plane may easily lead to unwanted motion blur. We therefore consider an event camera as a highly interesting and much more dynamic alternative visual sensor for this particular scenario.




Fig. 4. Left: *The Ackermann steering model with the ICR [14]. Both a left and a right turn are illustrated.* Right: *Connections between vehicle displacement, extrinsic transformation, and relative camera pose.*

4.1 Homographic Mapping and Bounding Box Extraction

We rely the globally-optimal BnB solver for correspondence-less AGV motion presented in [14], which also employs a normal, downward facing camera. We employ the two-dimensional Ackermann steering model describing the commonly non-holonomic motion of an AGV. Employing this 2-DoF model leads to benefits in BnB, the complexity of which strongly depends on the dimensionality of the solution space. As illustrated in Fig. 4, the Ackermann model constrains the motion of the vehicle to follow a circular-arc trajectory about an Instantaneous Centre of Rotation (ICR). The motion between successive frames can be conveniently described at the hand of two parameters: the half-angle of the relative rotation angle ϕ , and the baseline between the two views ρ . However, the alignment of the events requires a temporal parametrisation of the relative pose, which is why we employ the angular velocity $\omega = \frac{\theta}{t} = \frac{2\phi}{t}$ as well as the translational velocity $v = \omega r = \omega \rho \frac{1}{2 \sin(\phi)}$ in our model. The relative transformation from vehicle frame v' back to v is therefore given by

$$\mathbf{R}_v = \begin{bmatrix} \cos(\omega t) & -\sin(\omega t) & 0 \\ \sin(\omega t) & \cos(\omega t) & 0 \\ 0 & 0 & 1 \end{bmatrix} \text{ and } \mathbf{t}_v = \frac{v}{\omega} \begin{bmatrix} 1 - \cos(\omega t) \\ \sin(\omega t) \\ 0 \end{bmatrix}. \quad (8)$$

Further details about the derivation are given in the supplementary material.

In practice the vehicle frame hardly coincides with the camera frame. The orientation and the height of the origin can be chosen to be identical, and the camera may be laterally mounted in the centre of the vehicle. However, there is likely to be a displacement along the forward direction, which we denote by the signed variable s . In other words, $\mathbf{R}_v^c = \mathbf{I}_{3 \times 3}$ and $\mathbf{t}_v^c = [0 \ s \ 0]^T$. As illustrated in Fig. 4, the transformation from camera pose c' (at an arbitrary future timestamp) to c (at the initial timestamp t_{ref}) is therefore given by

$$\begin{aligned} \mathbf{R}_c &= \mathbf{R}_v^{cT} \mathbf{R}_v \mathbf{R}_v^c, \\ \mathbf{t}_c &= -\mathbf{R}_v^{cT} \mathbf{t}_v^c + \mathbf{R}_v^{cT} \mathbf{t}_v + \mathbf{R}_v^{cT} \mathbf{R}_v \mathbf{t}_v^c. \end{aligned} \quad (9)$$

Using the known plane normal vector $\mathbf{n} = [0 \ 0 \ -1]^T$ and depth-of-plane d , the image warping function $W(\mathbf{x}_k, t_k; [\omega \ v]^T)$ that permits the transfer of an event $e_k = \{\mathbf{x}_k, t_k, s_k\}$ into the reference view at t_{ref} is finally given by the planar homography equation

$$\mathbf{H} [\mathbf{x}_k^T \ 1]^T = \mathbf{K} (\mathbf{R}_c - \frac{\mathbf{t}_c \mathbf{n}^T}{d}) \mathbf{K}^{-1} [\mathbf{x}_k^T \ 1]^T. \quad (10)$$

Note that \mathbf{K} here denotes a regular perspective camera calibration matrix with homogeneous focal length f , zero skew, and a principal point at $[u_0 \ v_0]^T$. Note further that the substituted time parameter needs to be equal to $t = t_k - t_{\text{ref}}$, and that the result needs to be dehomogenised. After expansion, we easily obtain

$$\begin{aligned} \mathbf{x}'_k &= W(\mathbf{x}_k, t_k; [\omega \ v]^T) = [x'_k \ y'_k]^T \\ &= \begin{bmatrix} -[y_k - v_0 + s \frac{f}{d}] \sin(\omega t) + [x_k - u_0 - \frac{f}{d} (\frac{v}{w})] \cos(\omega t) + \frac{f}{d} (\frac{v}{w}) + u_0 \\ [x_k - u_0 - \frac{f}{d} (\frac{v}{w})] \sin(\omega t) + [y_k - v_0 + s \frac{f}{d}] \cos(\omega t) - s \frac{f}{d} + v_0 \end{bmatrix}. \end{aligned} \quad (11)$$

Finally, the bounding box \mathcal{P}_k^Θ is found by bounding the values of x'_k and y'_k over the intervals $\omega \in \mathcal{W} = [\omega_{\min}; \omega_{\max}]$ and $v \in \mathcal{V} = [v_{\min}; v_{\max}]$. The bounding is easily achieved if simply considering monotonicity of functions over given sub-branches. For example, if $\omega_{\min} \geq 0$, $v_{\min} \geq 0$, $x_k \geq u_0$, and $y_k \geq v_0 - s\frac{f}{d}$, we obtain

$$\begin{aligned} \underline{x}'_k &= -[y_k - v_0 + s\frac{f}{d}] \sin(\omega_{\max}t) + [x_k - u_0 - \frac{f}{d}(\frac{v_{\min}}{w_{\max}})] \cos(\omega_{\max}t) + \frac{f}{d}(\frac{v_{\min}}{w_{\max}}) + u_0, \\ \overline{x}'_k &= -[y_k - v_0 + s\frac{f}{d}] \sin(\omega_{\min}t) + [x_k - u_0 - \frac{f}{d}(\frac{v_{\max}}{w_{\min}})] \cos(\omega_{\min}t) + \frac{f}{d}(\frac{v_{\max}}{w_{\min}}) + u_0, \\ \underline{y}'_k &= [x_k - u_0 - \frac{f}{d}(\frac{v_{\max}}{w_{\min}})] \sin(\omega_{\min}t) + [y_k - v_0 + s\frac{f}{d}] \cos(\omega_{\max}t) - s\frac{f}{d} + v_0, \text{ and} \\ \overline{y}'_k &= [x_k - u_0 - \frac{f}{d}(\frac{v_{\min}}{w_{\max}})] \sin(\omega_{\max}t) + [y_k - v_0 + s\frac{f}{d}] \cos(\omega_{\min}t) - s\frac{f}{d} + v_0. \end{aligned} \quad (12)$$

We kindly refer the reader to the supplementary material for all further cases.

5 Experimental Evaluation

We present two suites of experiments. The first one validates the global optimality, accuracy and robustness of our solver on simulated data. The second one then applies it to the real-world scenario of AGV motion estimation.

5.1 Accuracy and Robustness of Globally Optimal Motion Estimation

We start by evaluating the accuracy of the motion estimation with contrast maximisation function L_{SoS} over synthetic data. As already implied in [11], L_{SoS} can be considered as a solid starting point for the evaluation. Our synthetic data consists of randomly generated horizontal and vertical line segments on a plane at a depth of 2.0 m. We consider Ackermann motion with an angular velocity $\omega = 28.6479^\circ/\text{s}$ (0.5 rad/s) and a linear velocity $v = 0.5 \text{ m/s}$. Events are generated by randomly choosing a 3D point on a line, and reprojecting it into a random camera pose sampled by a random timestamp within the interval $[0, 0.1 \text{ s}]$. The result of our method is finally evaluated by running BnB over the search space $\mathcal{W} = [0.4, 0.6]$ and $\mathcal{V} = [0.4, 0.6]$, and comparing the retrieved solution against the result of an exhaustive search with sampling points every $\delta\omega = 0.001 \text{ rad/s}$ and $\delta v = 0.001 \text{ m/s}$. BnB is furthermore configured to terminate the search if $|\omega_{\max} - \omega_{\min}| \leq 0.00078 \text{ rad/s}$ or $|v_{\max} - v_{\min}| \leq 0.00078 \text{ m/s}$. The experiment is repeated 1000 times.

Figures 5(a) and 5(b) illustrate the distribution of the errors for both methods in the noise-free case. The standard deviation of the exhaustive search and BnB are $\sigma_\omega = 1.0645^\circ/\text{s}$, $\sigma_v = 0.0151 \text{ m/s}$ and $\sigma_\omega = 1.305^\circ/\text{s}$, $\sigma_v = 0.0150 \text{ m/s}$, respectively. While this result suggests that BnB works well and sustainably returns a result very close to the optimum found by exhaustive search, we still note that the optimum identified by both methods has a bias with respect to ground truth, even in the noise-free case. Note however that this is related to




Fig. 5. Simulation results. (a) and (b) indicate the error distribution for ω and v over all experiments for both our proposed method as well as an exhaustive search. (c) and (d) visualise the average error of the estimated parameters caused by additional salt and pepper noise in the event stream. Results are averaged over 1000 random experiments. Note that our proposed method has excellent robustness even for N/E ratios up to 40%.

the nature of the contrast maximisation function, and not our globally optimal solution strategy.

In order to analyse robustness, we randomly add salt and pepper noise to the event stream with noise-to-event (N/E) ratios between 0 and 0.4 (Example objective functions for different N/E ratios have already been illustrated in Fig. 2). Figure 5(c) and 5(d) show the error for each noise level again averaged over 1000 experiments. As can be observed, the errors are very similar and behave more or less independently of the amount of added noise. The latter result underlines the high robustness of our approach.

5.2 Application to Real Data and Comparison Against Alternatives

We apply our method to real data collected by a DAVIS346 event camera, which outputs events streams with a maximum time resolution of $1\mu\text{s}$ as well as regular frames at a frame rate of 30 Hz. Images have a resolution of 346×260 . We mount the camera on the front of a XQ-4 Pro robot and let it face downward. The displacement from the non-steering axis to the camera is $s = -0.45$ m, and the height difference between camera and ground is $d = 0.23$ m. We recorded several motion sequences on a wood grain foam which has highly self-similar texture and poses a challenge to reliably extract and match features. Ground truth is obtained via an Optitrack optical motion tracking system. Our algorithm is working in undistorted coordinates, which is why normalisation and undistortion are computed in advance. The following aspects are evaluated:

Different objective functions: We test the algorithm with all aforementioned six contrast functions over various types of motions, including a straight line, a circle, and an arbitrarily curved trajectory. Table 3 shows the RMS errors of the estimated dynamic parameters, and compares the accuracy of all six alternatives. We furthermore apply two state-of-the-art approaches for regular images, namely the correspondence-less globally optimal feature-based approach (GOVO) from [14], as well as the Improved Fourier Mellin Invariant

transform (IFMI) in [4,31]. Even though these alternatives use the same non-holonomic or planar motion models, event-based motion estimation methods significantly outperform the intensity-camera-based alternatives (L_{SoSAaS} on top, and L_{SoS} and L_{Var} also have good performance).

Table 3. RMS errors for different datasets and methods.

Method	Line w [°/s]	Line v [m/s]	Circle w [°/s]	Circle v [m/s]	Curve w [°/s]	Curve v [m/s]
SoE	2.4089	0.0158	2.2121	0.0252	3.6282	0.0263
SoEaS	2.4057	0.0158	2.0178	0.0242	3.6282	0.0263
SoS	0.5127	0.0086	1.0884	0.0083	3.0091	0.0208
SoSA	1.9606	0.0287	4.2496	0.0734	9.2904	0.0727
SoSAaS	0.5175	0.0086	0.5294	0.0046	0.5546	0.0189
Var	0.5127	0.0086	1.0884	0.0083	3.0091	0.0208
IFMI	145.3741	1.0594	8.1092	0.0243	12.8047	0.0192
GOVO	6.9705	0.2409	4.5506	0.0642	9.8652	0.0590

Event-Based vs Frame-Based: GOVO [14] and IFMI [31] are frame-based algorithms specifically designed for planar AGV motion estimation under featureless conditions. Figure 1 shows an example frame of the wood grain foam texture, and Fig. 7 the results obtained for all methods. As can be observed, GOVO finds as little as three corner features for some of the images, thus making it difficult to accurately recover the vehicle displacement despite the globally-optimal correspondence-less nature of the algorithm. Both IFMI and GOVO occasionally lose tracking (especially for linear motion), which leaves our proposed globally-optimal event-based method using L_{SoSAaS} as the best method.




Fig. 6. Estimated trajectories by our method (SoS), gradient ascent with various initializations, and ground truth (gt). The table indicates the RMS errors for the best performing gradient ascent run and SoS.

BnB vs Gradient Ascent: We apply both gradient descent as well as BnB to the *Foam* dataset with curved motion. For the first temporal interval and the local search method, we vary the initial angular velocity ω and linear velocity v between -1 and 0.8 with steps of 0.2 (rad/s or m/s, respectively). For later

intervals, we use the previous local optimum. Figure 6 illustrates the estimated trajectories for all initial values, compared against ground truth and a BnB search using L_{SoS} . RMS errors are also indicated. As clearly shown, even the best initial guess eventually diverges under a local search strategy, thus leading to clearly inferior results compared to our globally optimal search.

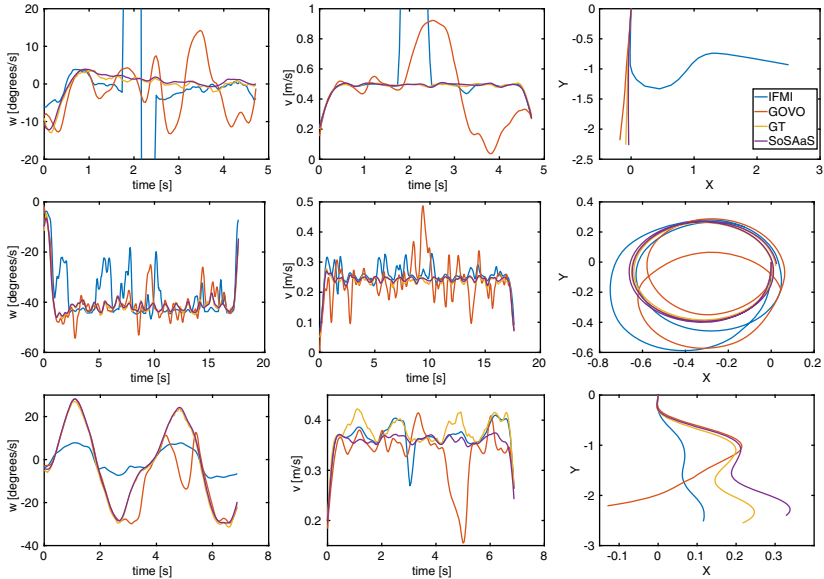


Fig. 7. Results for all methods over different datasets. The first two columns are errors over time for ω and v , and the third column illustrates a bird's eye view onto the integrated trajectories.

Various Textures: More results over datasets with other ground floor textures can be found in the supplementary material.

6 Discussion

We have introduced the first globally optimal solution to contrast maximisation for un-warped event streams. To the best of our knowledge, we are also the first to apply the idea of homography estimation via contrast maximisation to the real-world case of non-holonomic motion estimation with a downward facing camera mounted on an AGV. The challenging conditions in this scenario favorise dynamic vision sensors over regular frame-based cameras, a claim that is supported by our experimental results. The latter furthermore prove that global solutions are important and significantly outperform incremental local refinement. The recursive formulation of our bounds lets us find the global optimum

over event streams of 0.04 s within less than one minute, a respectable achievement given the typically low computational efficiency of BnB solvers.

Acknowledgments. The authors would like to thank the fundings sponsored by Natural Science Foundation of Shanghai (grant number: 19ZR1434000) and Natural Science Foundation of China (grant number: 61950410612).

References

1. Aqel, M.O., Marhaban, M.H., Saripan, M.I., Ismail, N.B.: Review of visual odometry: types, approaches, challenges, and applications. *SpringerPlus* **5**(1), 1897 (2016)
2. Birem, M., Kleihorst, R., El-Ghouti, N.: Visual odometry based on the fourier transform using a monocular ground-facing camera. *J. Real-Time Image Proc.* **14**(3), 637–646 (2018)
3. Brandli, C., Berner, R., Yang, M., Liu, S.C., Delbruck, T.: A 240×180 130 db 3 μ s latency global shutter spatiotemporal vision sensor. *IEEE J. Solid-State Circ.* **49**(10), 2333–2341 (2014)
4. Bülow, H., Birk, A.: Fast and robust photomapping with an unmanned aerial vehicle (UAV). In: 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3368–3373. IEEE (2009)
5. Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., Leonard, J.J.: Past, present, and future of simultaneous localization and mapping: toward the robust-perception age. *IEEE Trans. Rob.* **32**(6), 1309–1332 (2016)
6. Campbell, D., Petersson, L., Kneip, L., Li, H.: Globally-optimal inlier set maximization for simultaneous camera pose and feature correspondence. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1–10 (2017)
7. Chen, X., Vempati, A.S., Beardsley, P.: Streetmap-mapping and localization on ground planes using a downward facing camera. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 1672–1679. IEEE (2018)
8. Dille, M., Grocholsky, B., Singh, S.: Outdoor downward-facing optical flow odometry with commodity sensors. In: Howard, A., Iagnemma, K., Kelly, A. (eds.) *Field and Service Robotics*. Springer Tracts in Advanced Robotics, vol. 62, pp. 183–193. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13408-1_17
9. Fuentes-Pacheco, J., Ruiz-Ascencio, J., Rendón-Mancha, J.M.: Visual simultaneous localization and mapping: a survey. *Artif. Intell. Rev.* **43**(1), 55–81 (2015)
10. Gallego, G., et al.: Event-based vision: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* (2020)
11. Gallego, G., Gehrig, M., Scaramuzza, D.: Focus is all you need: loss functions for event-based vision. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 12280–12289 (2019)
12. Gallego, G., Rebecq, H., Scaramuzza, D.: A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3867–3876 (2018)
13. Gallego, G., Scaramuzza, D.: Accurate angular velocity estimation with an event camera. *IEEE Robot. Autom. Lett.* **2**(2), 632–639 (2017)
14. Gao, L., Su, J., Cui, J., Zeng, X., Peng, X., Kneip, L.: Efficient globally-optimal correspondence-less visual odometry for planar ground vehicles. In: 2020 International Conference on Robotics and Automation (ICRA), pp. 2696–2702. IEEE (2020)

15. Gonzalez, R., Rodriguez, F., Guzman, J.L., Pradalier, C., Siegwart, R.: Combined visual odometry and visual compass for off-road mobile robots localization. *Robotica* **30**(6), 865–878 (2012)
16. Guo, X., Xu, Z., Lu, Y., Pang, Y.: An application of Fourier-Mellin transform in image registration. In: *The Fifth International Conference on Computer and Information Technology (CIT 2005)*, pp. 619–623. IEEE (2005)
17. Huang, K., Wang, Y., Kneip, L.: Motion estimation of non-holonomic ground vehicles from a single. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 12706–12715. IEEE (2019)
18. Jordan, J., Zell, A.: Ground plane based visual odometry for RGBD-cameras using orthogonal projection. *IFAC-PapersOnLine* **49**(15), 108–113 (2016)
19. Kazik, T., Göktoğan, A.H.: Visual odometry based on the Fourier-Mellin transform for a rover using a monocular ground-facing camera. In: *2011 IEEE International Conference on Mechatronics*, pp. 469–474. IEEE (2011)
20. Lovegrove, S., Davison, A.J., Ibanez-Guzmán, J.: Accurate visual odometry from a rear parking camera. In: *2011 IEEE Intelligent Vehicles Symposium (IV)*, pp. 788–793. IEEE (2011)
21. Mitrokhin, A., Fermüller, C., Parameshwara, C., Aloimonos, Y.: Event-based moving object detection and tracking. In: *2018 IEEE RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1–9 (2018)
22. Nourani-Vatani, N., Borges, P.V.K.: Correlation-based visual odometry for ground vehicles. *J. Field Robot.* **28**(5), 742–768 (2011)
23. Nourani-Vatani, N., Roberts, J., Srinivasan, M.V.: Practical visual odometry for car-like vehicles. In: *2009 IEEE International Conference on Robotics and Automation*, pp. 3551–3557. IEEE (2009)
24. Peng, X., Cui, J., Kneip, L.: Articulated multi-perspective cameras and their application to truck motion estimation. In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2052–2059. IEEE (2019)
25. Piyathilaka, L., Munasinghe, R.: An experimental study on using visual odometry for short-run self localization of field robot. In: *2010 Fifth International Conference on Information and Automation for Sustainability*, pp. 150–155. IEEE (2010)
26. Rebecq, H., Gallego, G., Mueggler, E., Scaramuzza, D.: EMVS: event-based multi-view stereo-3D reconstruction with an event camera in real-time. *Int. J. Comput. Vis.* **126**(12), 1394–1414 (2018)
27. Stoffregen, T., Gallego, G., Drummond, T., Kleeman, L., Scaramuzza, D.: Event-based motion segmentation by motion compensation. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 7244–7253 (2019)
28. Stoffregen, T., Kleeman, L.: Simultaneous optical flow and segmentation (SOFAS) using dynamic vision sensor. In: *2017 Australasian Conference on Robotics and Automation (ACRA)*, pp. 52–61 (2017)
29. Stoffregen, T., Kleeman, L.: Event cameras, contrast maximization and reward functions: an analysis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 12300–12308 (2019)
30. Wang, Y., Huang, K., Peng, X., Li, H., Kneip, L.: Reliable frame-to-frame motion estimation for vehicle-mounted surround-view camera systems. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1660–1666. IEEE (2020)
31. Xu, Q., Chavez, A.G., Bülow, H., Birk, A., Schwertfeger, S.: Improved Fourier Mellin invariant for robust rotation estimation with omni-cameras. In: *2019 IEEE International Conference on Image Processing (ICIP)*, pp. 320–324. IEEE (2019)

32. Ye, C., Mitrokhin, A., Parameshwara, C., Fermüller, C., Yorke, J.A., Aloimonos, Y.: Unsupervised learning of dense optical flow and depth from sparse event data. CoRR abs/1809.08625 (2018). <http://arxiv.org/abs/1809.08625>
33. Yu, Y., Pradalier, C., Zong, G.: Appearance-based monocular visual odometry for ground vehicles. In: 2011 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM), pp. 862–867. IEEE (2011)
34. Zhu, A.Z., Atanasov, N., Daniilidis, K.: Event-based feature tracking with probabilistic data association. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), pp. 4465–4470. IEEE (2017)
35. Zhu, A.Z., Chen, Y., Daniilidis, K.: Realtime time synchronized event-based stereo. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11210, pp. 438–452. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01231-1_27
36. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: EV-FlowNet: self-supervised optical flow estimation for event-based cameras. arXiv preprint [arXiv:1802.06898](https://arxiv.org/abs/1802.06898) (2018)
37. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: Unsupervised event-based learning of optical flow, depth, and egomotion. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 989–997 (2019)
38. Zienkiewicz, J., Davison, A.: Extrinsic autocalibration for dense planar visual odometry. *J. Field Robot.* **32**(5), 803–825 (2015)