

Secrets of Event-Based Optical Flow

Shintaro Shiba^{1,2}, Yoshimitsu Aoki¹, and Guillermo Gallego^{2,3}

¹ Department of Electronics and Electrical Engineering, Faculty of Science and Technology, Keio University, Kanagawa, Japan. sshiba@keio.jp

² Department of EECS, Technische Universität Berlin, Berlin, Germany

³ Einstein Center Digital Future and SCIOI Excellence Cluster, Berlin, Germany

Abstract. Event cameras respond to scene dynamics and offer advantages to estimate motion. Following recent image-based deep-learning achievements, optical flow estimation methods for event cameras have rushed to combine those image-based methods with event data. However, it requires several adaptations (data conversion, loss function, etc.) as they have very different properties. We develop a principled method to extend the Contrast Maximization framework to estimate optical flow from events alone. We investigate key elements: how to design the objective function to prevent overfitting, how to warp events to deal better with occlusions, and how to improve convergence with multi-scale raw events. With these key elements, our method ranks first among unsupervised methods on the MVSEC benchmark, and is competitive on the DSEC benchmark. Moreover, our method allows us to expose the issues of the ground truth flow in those benchmarks, and produces remarkable results when it is transferred to unsupervised learning settings. Our code is available at https://github.com/tub-rip/event_based_optical_flow

1 Introduction

Event cameras are novel bio-inspired vision sensors that naturally respond to motion of edges in image space with high dynamic range (HDR) and minimal blur at high temporal resolution (on the order of μs) [3, 35]. These advantages provide a rich signal for accurate motion estimation in difficult real-world scenarios for frame-based cameras. However such a signal is, by nature, asynchronous and sparse, which is not compatible with traditional computer vision algorithms. This poses the challenge of rethinking visual processing [12, 23]: motion patterns (i.e., *optical flow*) are no longer obtained by analyzing the intensities of images captured at regular intervals, but by analyzing the stream of events (per-pixel brightness changes) produced by the event camera.

Multiple methods have been proposed for event-based optical flow estimation. They can be broadly categorized in two: (i) model-based methods, which investigate the principles and characteristics of event data that enable optical flow estimation, and (ii) learning-based methods, which exploit correlations in the data and/or apply the above-mentioned principles to compute optical flow. One of the challenges of event-based optical flow is the lack of ground truth flow in real-world datasets (at μs resolution and HDR) [12], which makes it difficult

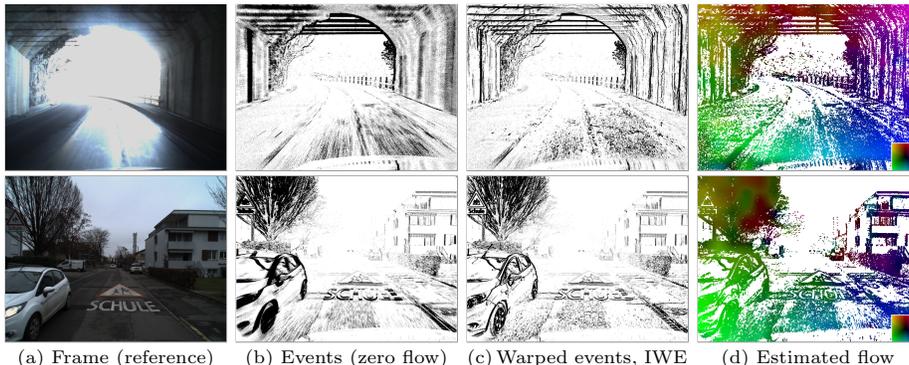


Fig. 1. Two test sequences (interlaken_00.b, thun_01.a) from the DSEC dataset [17]. Our optical flow estimation method produces sharp images of warped events (IWE) despite the scene complexity, the large pixel displacement and the high dynamic range. The examples utilize 500k events on an event camera with 640×480 pixels.

to evaluate and compare the methods properly, and to train supervised learning-based ones. Ground truth (GT) in de facto standard datasets [17, 55] is given by the *motion field* [49] using additional depth sensors and camera information. However, such data is limited by the field-of-view (FOV) and resolution (spatial and temporal) of the depth sensor, which do not match those of event cameras. Hence, it is paramount to develop interpretable optical flow methods that exploit the characteristics of event data, and that do not need expensive-to-collect and error-prone ground truth.

Among prior work, Contrast Maximization (CM) [13, 14] is a powerful framework that allows us to tackle multiple motion estimation problems (rotational motion [15, 20, 22], homographic motion [14, 30, 34], feature flow estimation [39, 44, 53, 54], motion segmentation [27, 32, 42, 52], and also reconstruction [14, 36, 51]). It maximizes an objective function (e.g., contrast) that measures the alignment of events caused by the same scene edge. The intuitive interpretation is to estimate the motion by recovering the sharp (motion-compensated) image of edge patterns that caused the events. Preliminary work on applying CM to estimate optical flow has reported a problem of overfitting to the data, producing undesired flows that warp events to few pixels or lines [57] (i.e., event collapse [41]). This issue has been tackled by changing the objective function, from contrast to the energy of an average timestamp image [21, 33, 57], but this loss is not straightforward to interpret and makes training difficult to converge [13].

Given the state-of-the-art performance of CM in low-DOF motion problems and its issues in more complex motions (dense flow), we think prior work may have rushed to use CM in unsupervised learning of complex motions. There is a gap in understanding how CM can be sensibly extended to estimate dense optical flow accurately. In this paper we fill this gap and learn a few “secrets” that are also applicable to overcome the issues of previous approaches.

We propose to extend CM for dense optical flow estimation via a tile-based approach covering the image plane. We present several distinctive contributions:

1. A *multi-reference* focus loss function to improve accuracy and discourage overfitting (Sec. 3.2).
2. A principled *time-aware flow* to better handle occlusions, formulating event-based optical flow as a transport problem via differential equations (Sec. 3.3).
3. A *multi-scale* approach on the raw events to improve convergence to the solution and avoid getting trapped in local optima (Sec. 3.4).

The results of our experimental evaluation are surprising: the above design choices are key to our simple, model-based tile-based method (Fig. 1) achieving the best accuracy among all state-of-the-art methods, including supervised-learning ones, on the de facto benchmark of MVSEC indoor sequences [56]. Since our method is interpretable and produces better event alignment than the ground truth flow, both qualitatively and quantitatively, the experiments also expose the limitations of the current “ground truth”. Finally, experiments demonstrate that the above key choices are transferable to unsupervised learning methods, thus guiding future design and understanding of more proficient Artificial Neural Networks (ANNs) for event-based optical flow estimation.

Because of the above, we believe that the proposed design choices deserve to be called “secrets” [46]. To the best of our knowledge, they are novel in the context of event-based optical flow estimation, e.g., no prior work considers constant flow along its characteristic lines, designs the multi-reference focus loss to tackle overfitting, or has explicitly defined multi-scale (i.e., multi-resolution) contrast maximization on the raw events.

2 Prior Work on Event-based Optical Flow Estimation

Given the identified advantages of event cameras to estimate optical flow, extensive research on this topic has been carried out. Prior work has proposed adaptations of frame-based approaches (block matching [26], Lucas-Kanade [4]), filter-banks [6,31], spatio-temporal plane-fitting [1,3], time surface matching [29], variational optimization on voxelized events [2], and feature-based contrast maximization [14,43,53]. For a detailed survey, we refer to [12].

Current state-of-the-art approaches are ANNs [10,18,21,24,56,57], largely inspired by frame-based optical flow architectures [37,48]. Non-spiking-based approaches need to additionally adapt the input signal, converting the events into a tensor representation (event frames, time surfaces, voxel grids, etc. [7,16]). These learning-based methods can be classified into supervised, semi-supervised or unsupervised (to be used in Tab. 1). In terms of architectures, the three most common ones are U-Net [56], FireNet [57] and RAFT [18].

Supervised methods train ANNs in simulation and/or real-data [16,18,45]. This requires accurate ground truth flow that matches the space-time resolution of event cameras. While this is no problem in simulation, it incurs in a performance gap when trained models are used to predict flow on real data [45]. Besides, real-world datasets have issues in providing accurate ground truth flow.

Semi-supervised methods use the grayscale images from a colocated camera (e.g., DAVIS [47]) as a supervisory signal: images are warped using the flow predicted by the ANN and their photometric consistency is used as loss function [10, 24, 56]. While such supervisory signal is easier to obtain than real-world ground truth flow, it may suffer from the limitations of frame-based cameras (e.g., low dynamic range and motion blur), consequently affecting the trained ANNs. These approaches were pioneered by EV-FlowNet [56].

Unsupervised methods rely solely on event data. Their loss function consists of an event alignment error using the flow predicted by the ANN [21, 33, 50, 57]. Zhu et al. [57] extended EV-FlowNet [56] to the unsupervised setting using a motion-compensation loss function in terms of average timestamp images. This approach has been used and improved in [21, 33]. Paredes-Vallés et al. [33] also proposed FireFlowNet, a lightweight ANN producing competitive results.

Similar to the works in the unsupervised category, our method produces dense optical flow and does not need ground truth or additional supervisory signals. In contrast to prior work, we adopt a more classical modeling perspective to gain insight into the problem and discover principled solutions that can subsequently be applied to the learning-based setting. Stemming from an accurate and spatially-dependent contrast loss (the gradient magnitude [13]), we model the problem using a tile of patches and propose solutions to several problems: overfitting, occlusions and convergence. To the best of our knowledge, (i) no prior work has proposed to estimate dense optical flow from a CM model-based perspective, and (ii) no prior unsupervised learning approach based on motion compensation has succeeded in estimating optical flow without the average timestamp image loss.

3 Method

3.1 Event Cameras and Contrast Maximization

Event cameras have independent pixels that operate continuously and generate “events” $e_k \doteq (\mathbf{x}_k, t_k, p_k)$ whenever the logarithmic brightness at the pixel increases or decreases by a predefined amount, called contrast sensitivity. Each event e_k contains the pixel-time coordinates (\mathbf{x}_k, t_k) of the brightness change and its polarity $p_k = \{+1, -1\}$. Events occur asynchronously and sparsely on the pixel lattice, with a variable rate that depends on the scene dynamics.

The CM framework [14] assumes events $\mathcal{E} \doteq \{e_k\}_{k=1}^{N_e}$ are generated by moving edges, and transforms them geometrically according to a motion model \mathbf{W} , producing a set of warped events $\mathcal{E}'_{t_{\text{ref}}} \doteq \{e'_k\}_{k=1}^{N_e}$ at a reference time t_{ref} :

$$e_k \doteq (\mathbf{x}_k, t_k, p_k) \mapsto e'_k \doteq (\mathbf{x}'_k, t_{\text{ref}}, p_k). \quad (1)$$

The warp $\mathbf{x}'_k = \mathbf{W}(\mathbf{x}_k, t_k; \boldsymbol{\theta})$ transports each event from t_k to t_{ref} along the motion curve that passes through it. The vector $\boldsymbol{\theta}$ parametrizes the motion curves. Transformed events are aggregated on an image of warped events (IWE):

$$I(\mathbf{x}; \mathcal{E}'_{t_{\text{ref}}}, \boldsymbol{\theta}) \doteq \sum_{k=1}^{N_e} \delta(\mathbf{x} - \mathbf{x}'_k), \quad (2)$$

where each pixel \mathbf{x} sums the number of warped events \mathbf{x}'_k that fall within it. The Dirac delta δ is approximated by a Gaussian, $\delta(\mathbf{x} - \boldsymbol{\mu}) \approx \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \epsilon^2 \mathbf{Id})$ with $\epsilon = 1$ pixel. Next, an objective function $f(\boldsymbol{\theta})$ is built from the transformed events, such as the contrast of the IWE (2), given by the variance

$$\text{Var}(I(\mathbf{x}; \boldsymbol{\theta})) \doteq \frac{1}{|\Omega|} \int_{\Omega} (I(\mathbf{x}; \boldsymbol{\theta}) - \mu_I)^2 d\mathbf{x}, \quad (3)$$

with mean $\mu_I \doteq \frac{1}{|\Omega|} \int_{\Omega} I(\mathbf{x}; \boldsymbol{\theta}) d\mathbf{x}$. The objective function measures the goodness of fit between the events and the candidate motion curves (warp). Finally, an optimization algorithm iterates the above steps until convergence. The goal is to find the motion parameters that maximize the alignment of events caused by the same scene edge. Event alignment is measured by the strength of the edges of the IWE, which is directly related to image contrast [19].

Dense optical flow. In the task of interest, the warp used is [21, 33, 57]:

$$\mathbf{x}'_k = \mathbf{x}_k + (t_k - t_{\text{ref}}) \mathbf{v}(\mathbf{x}_k), \quad (4)$$

where $\boldsymbol{\theta} = \{\mathbf{v}(\mathbf{x})\}_{\mathbf{x} \in \Omega}$ is a flow field on the image plane at a set time, e.g., t_{ref} .

3.2 Multi-reference Focus Objective Function

Zhu et al. [57] report that the contrast objective (variance) overfits to the events. This is in part because the warp (4) can describe very complex flow fields, which can push the events to accumulate in few pixels [41]. To mitigate overfitting, we reduce the complexity of the flow field by dividing the image plane into a tile of non-overlapping patches, defining a flow vector at the center of each patch and interpolating the flow on all other pixels (we show the tiles in Sec. 3.4).

However, this is not enough. Additionally, we discover that warps that produce sharp IWEs *at any* reference time t_{ref} have a regularizing effect on the flow field, discouraging overfitting. This is illustrated in Fig. 2. In practice we compute the *multi-reference* focus loss using 3 reference times: t_1 (min), $t_{\text{mid}} \doteq (t_1 + t_{N_e})/2$ (midpoint) and t_{N_e} (max). The flow field is defined only at one reference time.

Furthermore, we measure event alignment using the magnitude of the IWE gradient because: (i) it has top accuracy performance among the objectives in [13], (ii) it is sensitive to the arrangement (i.e., permutation) of the IWE pixel values, whereas the variance of the IWE (3) is not, (iii) it converges more easily than other objectives we tested, (iv) it differs from the Flow Warp Loss (FWL) [45], which is defined using the variance (3) and will be used for evaluation.

Finally, letting the (squared) gradient magnitude of the IWE be

$$G(\boldsymbol{\theta}; t_{\text{ref}}) \doteq \frac{1}{|\Omega|} \int_{\Omega} \|\nabla I(\mathbf{x}; t_{\text{ref}})\|^2 d\mathbf{x}, \quad (5)$$

the proposed multi-reference focus objective function becomes the average of the G functions of the IWEs at multiple reference times:

$$f(\boldsymbol{\theta}) \doteq (G(\boldsymbol{\theta}; t_1) + 2G(\boldsymbol{\theta}; t_{\text{mid}}) + G(\boldsymbol{\theta}; t_{N_e})) / 4G(\mathbf{0}; -), \quad (6)$$

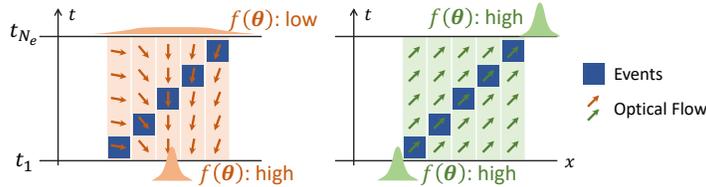


Fig. 2. *Multi-reference focus loss.* Assume an edge moves from left to right. Flow estimation with single reference time (t_1) can overfit to the data, warping all events into a single pixel, which results in a maximum contrast (at t_1). However, the same flow would produce low contrast (i.e., a blurry image) if events were warped to time t_{N_e} . Instead, we favor flow fields that produce high contrast (i.e., sharp images) at any reference time (here, $t_{\text{ref}} = t_1$ and $t_{\text{ref}} = t_{N_e}$). See results in Fig. 7.

normalized by the value of the G function with zero flow (identity warp). The normalization in (6) provides the same interpretation as the FWL: $f < 1$ implies the flow is worse than the zero flow baseline, whereas $f > 1$ means that the flow produces sharper IWEs than the baseline.

Remark: Warping to two reference times (min and max) was proposed in [57], but with important differences: (i) it was done for the average timestamp loss, hence it did not consider the effect on contrast or focus functions [13], and (ii) it had a completely different motivation: to lessen a back-propagation scaling problem, so that the gradients of the loss would not favor events far from t_{ref} .

3.3 Time-aware Flow

State-of-the-art event-based optical flow approaches are based on frame-based ones, and so they use the warp (4), which defines the flow $\mathbf{v}(\mathbf{x})$ as a function of \mathbf{x} (i.e., a pixel displacement between two given frames). However, this does not take into account the space-time nature of events, which is the basis of the CM approach, because not all events at a pixel \mathbf{x}_0 are triggered at the same timestamp t_k . They do not need to be warped with the same velocity $\mathbf{v}(\mathbf{x}_0)$. Figure 3 illustrates this with an occlusion example taken from the `slider_depth` sequence [28]. Instead of $\mathbf{v}(\mathbf{x})$, the *event-based flow* should be a function of space-time, $\mathbf{v}(\mathbf{x}, t)$, i.e. *time-aware*, and each event e_k should be warped according to the flow defined at (\mathbf{x}_k, t_k) . Let us propose a more principled warp than (4).

To define a space-time flow $\mathbf{v}(\mathbf{x}, t)$ that is compatible with the propagation of events along motion curves, we are inspired by the method of characteristics [11]. Just like the brightness constancy assumption states that brightness is constant along the true motion curves in image space, we assume the flow is constant along its streamlines: $\mathbf{v}(\mathbf{x}(t), t) = \text{const}$ (Fig. 3). Differentiating in time and applying the chain rule gives a system of partial differential equations (PDEs):

$$\frac{\partial \mathbf{v}}{\partial \mathbf{x}} \frac{d\mathbf{x}}{dt} + \frac{\partial \mathbf{v}}{\partial t} = \mathbf{0}, \quad (7)$$

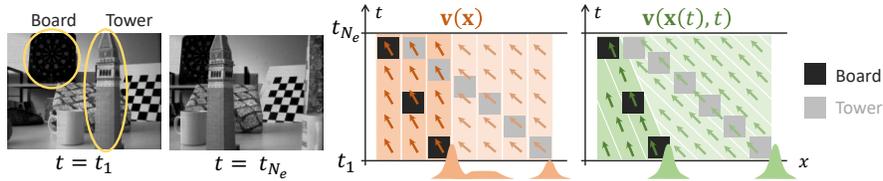


Fig. 3. *Time-aware Flow.* Traditional flow (4), inherited from frame-based approaches, assumes per-pixel constant flow $\mathbf{v}(\mathbf{x}) = \text{const}$, which cannot handle occlusions properly. The proposed space-time flow assumes constancy along streamlines, $\mathbf{v}(\mathbf{x}(t), t) = \text{const}$, which allows us to handle occlusions more accurately. (See results in Fig. 8)

where, as usual, $\mathbf{v} = d\mathbf{x}/dt$ is the flow. The boundary condition is given by the flow at say $t = 0$: $\mathbf{v}(\mathbf{x}, 0) = \mathbf{v}^0(\mathbf{x})$. This system of PDEs essentially states how to propagate (i.e., *transport*) a given flow $\mathbf{v}^0(\mathbf{x})$, from the boundary $t = 0$ to the rest of space \mathbf{x} and time t . The PDEs have advection terms and others that resemble those of the inviscid Burgers' equation [11] since the flow is transporting itself. We parametrize the flow at $t = t_{\text{mid}}$ (boundary condition), and then propagate it to the volume that encloses the current set of events \mathcal{E} . We develop two explicit methods to solve the PDEs, one with upwind differences and one with a conservative scheme adapted to Burgers' terms [40]. Each event e_k is then warped according to a flow $\hat{\mathbf{v}}$ given by the solution of the PDEs at (\mathbf{x}_k, t_k) :

$$\mathbf{x}'_k = \mathbf{x}_k + (t_k - t_{\text{ref}}) \hat{\mathbf{v}}(\mathbf{x}_k, t_k). \quad (8)$$

3.4 Multi-scale Approach

Inspired by classical estimation methods, we combine our tile-based approach with a multi-scale strategy. The goal is to improve the convergence of the optimizer in terms of speed and robustness (i.e., avoiding local optima).

Some learning-based works [33, 56, 57] also have a multi-scale component, inherited from the use of a U-Net architecture. However, they work on discretized event representations (voxel grid, etc.) to be compatible with CNNs. In contrast, our tile-based approach works directly on raw events, without discarding or quantizing the temporal information in the event stream. While some prior work outside the context of optical flow has considered multi-resolution on raw events [25], there is no agreement on the best way to perform multi-resolution due to the sparse and asynchronous nature of events.

Our multi-scale CM approach is illustrated in Fig. 4. For an event set \mathcal{E}_i , we apply the tile-based CM in a coarse-to-fine manner (e.g., $N_\ell = 5$ scales). There are $2^{l-1} \times 2^{l-1}$ tiles at the l -th scale. We use bilinear interpolation to upscale between any two scales. If there is a subsequent set \mathcal{E}_{i+1} , the flow estimated from \mathcal{E}_i is used to initialize the flow for \mathcal{E}_{i+1} . This is done by downsampling the finest flow to coarser scales. The coarsest scale initializes the flow for \mathcal{E}_{i+1} . For finer scales, initialization is computed as the average of the upsampled flow from the coarser scale of \mathcal{E}_{i+1} and the same-scale flow from \mathcal{E}_i .

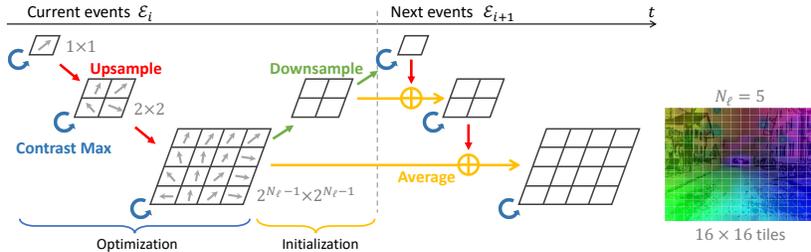


Fig. 4. Multi-scale Approach using tiles (rectangles) and raw events.

Composite Objective. To encourage additional smoothness of the flow, even in regions with few events, we include a flow regularizer $\mathcal{R}(\theta)$. The flow is obtained as the solution to the problem with the composite objective:

$$\theta^* = \arg \min_{\theta} (1/f(\theta) + \lambda \mathcal{R}(\theta)), \quad (9)$$

where, $\lambda > 0$ is the regularizer weight, and we use the total variation (TV) [38] as regularizer. We choose $1/f$ instead of simply $-f$ because it is convenient for ANN training, as we will apply in Sec 4.6.

4 Experiments

4.1 Datasets, Metrics and Hyper-parameters

We evaluate our method on sequences from the MVSEC dataset [55, 56], which is the de facto dataset used by prior works to benchmark optical flow. It provides events, grayscale frames, IMU data, camera poses, and scene depth from a LiDAR [55]. The dataset was extended in [56] to provide ground truth optical flow, computed as the motion field [49] given the camera velocity and the depth of the scene. The event camera has 346×260 pixel resolution [47]. In total, we evaluate on 63.5 million events spanning 265 seconds.

We also evaluate on a recent dataset that provides ground truth flow: DSEC [18]. It consists of sequences recorded with Prophesee Gen3 event cameras, of higher resolution (640×480 pixels), mounted on a car. Optical flow is also computed as the motion field, with the scene depth from a LiDAR. In total, we evaluate on 3 billion events spanning the 208 seconds of the test sequences.

The metrics used to assess optical flow accuracy are the average endpoint error (AEE) and the percentage of pixels with AEE greater than 3 pixels (denoted by “% Out”), both are measured over pixels with valid ground-truth and at least one event in the evaluation intervals. We also use the FWL metric (the IWE variance relative to that of the identity warp) to assess event alignment [45].

In all experiments our method uses $N_\ell = 5$ resolution scales, $\lambda = 0.0025$ in (9), and the Newton-CG optimization algorithm with a maximum of 20 iterations/scale. The flow at t_{mid} is transported to each side via the upwind or

Table 1. Results on MVSEC dataset [56]. Methods are sorted according to how much data they need: supervised learning (SL) requires ground truth flow; semi-supervised learning (SSL) uses grayscale images for supervision; unsupervised learning (USL) uses only events; and model-based (MB) needs no training data. Bold is the best among all methods; underlined is second best. Nagata et al. [29] evaluate on shorter time intervals; for comparison, we scale the errors to $dt = 1$.

		indoor_flying1		indoor_flying2		indoor_flying3		outdoor_day1	
		AEE ↓	%Out ↓	AEE ↓	%Out ↓	AEE ↓	%Out ↓	AEE ↓	%Out ↓
$dt = 1$									
SL	EV-FlowNet-EST [16]	0.97	0.91	1.38	8.20	1.43	6.47	-	-
	EV-FlowNet+ [45]	0.56	1.00	<u>0.66</u>	<u>1.00</u>	<u>0.59</u>	1.00	0.68	0.99
	E-RAFT [18]	-	-	-	-	-	-	0.24	1.70
SSL	EV-FlowNet (original) [56]	1.03	2.20	1.72	15.10	1.53	11.90	0.49	0.20
	Spike-FlowNet [24]	0.84	-	1.28	-	1.11	-	0.49	-
	Ziluo et al. [10]	0.57	0.10	0.79	1.60	0.72	1.30	0.42	0.00
USL	EV-FlowNet [57]	0.58	0.00	1.02	4.00	0.87	3.00	0.32	0.00
	EV-FlowNet (retrained) [33]	0.79	1.20	1.40	10.90	1.18	7.40	0.92	5.40
	FireFlowNet [33]	0.97	2.60	1.67	15.30	1.43	11.00	1.06	6.60
	ConvGRU-EV-FlowNet [21]	0.60	0.51	1.17	8.06	0.93	5.64	0.47	0.25
MB	Nagata et al. [29]	0.62	-	0.93	-	0.84	-	0.77	-
	Akolkar et al. [1]	1.52	-	1.59	-	1.89	-	2.75	-
	Brebion et al. [5]	<u>0.52</u>	0.10	0.98	5.50	0.71	2.10	0.53	0.20
	Ours (w/o time aware)	0.42	<u>0.09</u>	0.60	0.59	0.50	<u>0.29</u>	<u>0.30</u>	0.11
	Ours (Upwind)	0.42	0.10	0.60	0.59	0.50	0.28	<u>0.30</u>	<u>0.10</u>
	Ours (Burgers')	0.42	0.10	0.60	0.59	0.50	0.28	<u>0.30</u>	<u>0.10</u>
$dt = 4$									
SSL	EV-FlowNet (original) [56]	2.25	24.70	4.05	45.30	3.45	39.70	1.23	<u>7.30</u>
	Spike-FlowNet [24]	2.24	-	3.83	-	3.18	-	<u>1.09</u>	-
	Ziluo et al. [10]	1.77	14.70	<u>2.52</u>	26.10	<u>2.23</u>	22.10	0.99	3.90
USL	EV-FlowNet [57]	2.18	24.20	3.85	46.80	3.18	47.80	1.30	9.70
	ConvGRU-EV-FlowNet [21]	2.16	21.51	3.90	40.72	3.00	29.60	1.69	12.50
MB	Ours (w/o time aware)	1.68	12.79	2.49	<u>26.31</u>	2.06	18.93	1.25	9.19
	Ours (Upwind)	<u>1.69</u>	<u>12.83</u>	2.49	26.37	2.06	<u>19.02</u>	1.25	9.23
	Ours (Burgers')	<u>1.69</u>	12.95	2.49	26.35	2.06	19.03	1.25	9.21
	Ours (Upwind)	<u>1.69</u>	12.95	2.49	26.35	2.06	19.03	1.25	9.21

Burgers' PDE solver (using 5 bins for MVSEC, 40 for DSEC), and used for event warping (8) (see Suppl. Mat.). In the optimization, we use 30k events for MVSEC indoor sequences, 40k events for outdoors, and 1.5M events for DSEC.

4.2 Results on MVSEC

Table 1 reports the results on the MVSEC benchmark. The different methods (rows) are compared on three indoor sequences and one outdoor sequence (columns). This is because many learning-based methods train on the other outdoor sequence, which is therefore not used for testing. Following Zhu et al., outdoor_day1 is tested only on specified 800 frames [56]. The top part of Tab. 1 reports the flow corresponding to a time interval of $dt = 1$ grayscale frame (at $\approx 45\text{Hz}$, i.e., 22.2ms), and the bottom part corresponds to $dt = 4$ frames (89ms).

Our methods provide the best results among all methods in all indoor sequences and are the best among the unsupervised and model-based methods in the outdoor sequence. The errors for $dt = 4$ are about four times larger than

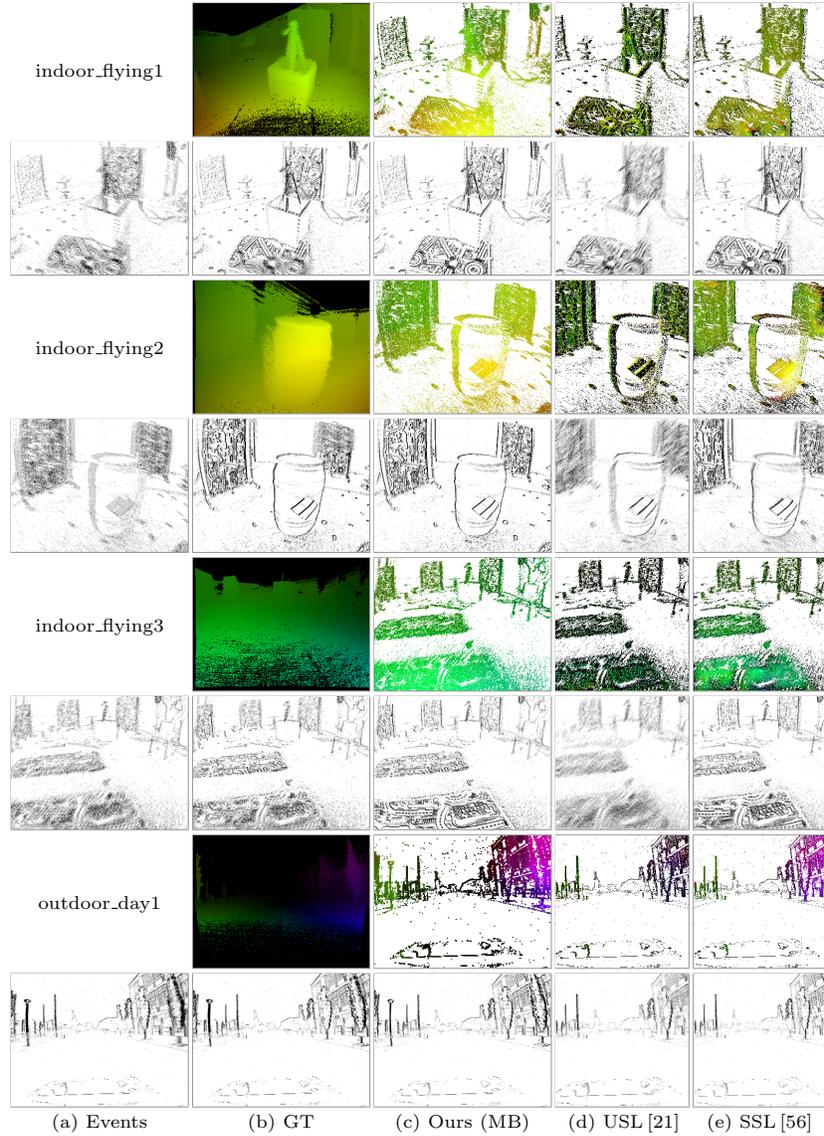


Fig. 5. *MVSEC comparison* ($dt = 4$) of our method and two state-of-the-art baselines: ConvGRU-EV-FlowNet (USL) [21] and EV-FlowNet (SSL) [56]. For each sequence, the upper row shows the flow masked by the input events, and the lower row shows the IWE using the flow. Our method produces the sharpest motion-compensated IWEs. Note that learning-based methods crop input events to center 256×256 pixels, whereas our method does not. Black points in ground truth (GT) flow maps indicate the absence of LiDAR measurements. The optical flow color wheel is in Fig. 1.

Table 2. Results on DSEC test sequences [18]. For the calculation of FWL, we use events within 100ms. More sequences are provided in the supplementary material.

	thun_01_a			thun_01_b			zurich_city_15_a		
	AEE ↓	%Out ↓	FWL ↑	AEE ↓	%Out ↓	FWL ↑	AEE ↓	%Out ↓	FWL ↑
E-RAFT [18]	0.65	1.87	1.20	0.58	1.52	1.18	0.59	1.30	1.34
Ours	2.12	17.68	1.24	2.48	23.56	1.24	2.35	20.99	1.41

those for $dt = 1$, which is sensible given the ratio of time interval sizes. We observe no significant differences between the three versions of the method tested (warp models, see also Sec. 4.5), which can be attributed to the fact that the MVSEC dataset does not comprise large pixel displacements or occlusions.

Qualitative results are shown in Fig. 5, where we compare our method against the state of the art. Our method provides sharper IWEs than the baselines, without overfitting, and the estimated flow resembles the ground truth one.

Ground truth (GT) is not available on the entire image plane (see Fig. 5), such as in pixels not covered by the LiDAR’s range, FOV, or spatial sampling. Additionally, there may be interpolation issues in the GT, since the LiDAR works at 20 Hz and the GT flow is given at frame rate (45 Hz). In the outdoor sequences, the GT from the LiDAR and the camera motion cannot provide correct flow for independently moving objects (IMOs). These issues of the GT are noticeable in the IWEs: they are not as sharp as expected. In contrast, the IWEs produced by our method are sharp. Taking now into account the GT quality on the comparison Table 1, it is remarkable that our method outperforms the state-of-the-art baselines on the indoor sequences, where GT has the best quality (with more points in the valid LiDAR range and no IMOs).

4.3 Results on DSEC

Table 2 gives quantitative results on the DSEC Optical Flow benchmark. Currently only the method that proposed the benchmark reports values [18]. As expected, this supervised learning method is better than ours in terms of flow accuracy because (i) it has additional training information (GT labels), and (ii) it is trained using the same type of GT signal used in the evaluation. Nevertheless, our method provides competitive results and is better in terms of FWL, which exposes similar GT quality issues as those of MVSEC: pixels without GT (LiDAR’s FOV and IMOs). Qualitative results are shown in Fig. 6. Our method provides sharp IWEs, even for IMOs (car) and the road close to the camera.

We observe that the evaluation intervals (100ms) are large for optical flow standards. In the benchmark, 80% of the GT flow has up to 22px displacement, which means that 20% of the GT flow is larger than 22px (on VGA resolution). The apparent motion during such intervals is sufficiently large that it breaks the classical assumption of scene points flowing in linear trajectories.

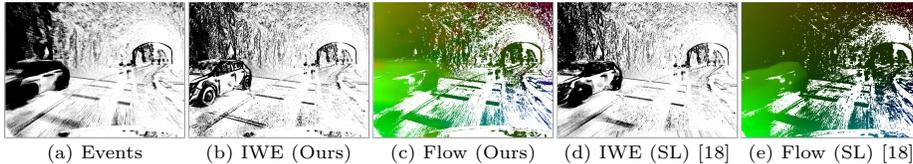


Fig. 6. *DSEC results* on the interlaken_00_b test sequence (no GT available). Since GT is missing at IMOs and points outside the LiDAR’s FOV, the supervised method [18] may provide inaccurate predictions around IMOs and road points close to the camera, whereas our method produces sharp edges. For visualization, we use 1M events.

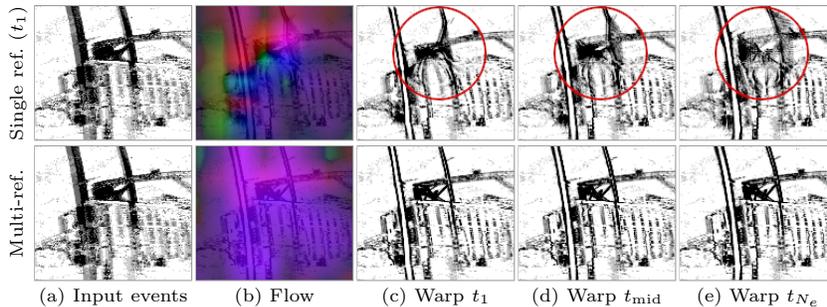


Fig. 7. *Effect of the multi-reference focus loss.*

4.4 Effect of the Multi-reference Focus Loss

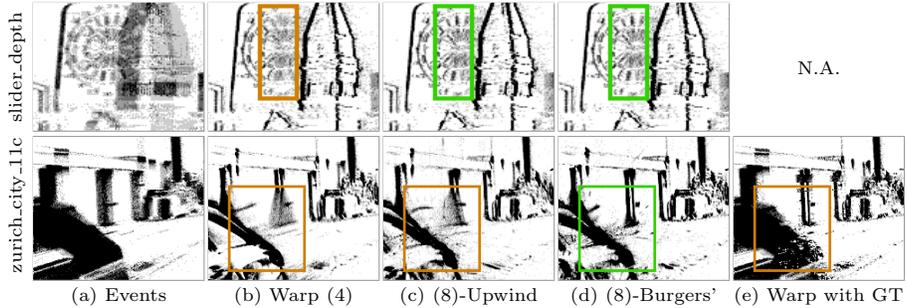
The effect of the proposed multi-reference focus loss is shown in Fig. 7. The single-reference focus loss function can easily overfit to the only reference time, pushing all events into a small region of the image at t_1 while producing blurry IWEs at other times (t_{mid} and t_{N_e}). Instead, our proposed multi-reference focus loss discourages such overfitting, as the loss favors flow fields which produce sharp IWEs at *any* reference time. The difference is also noticeable in the flow: the flow from the single-reference loss is irregular, with a lot of spatial variability in terms of directions (many colors, often in opposite directions of the color wheel). In contrast, the flow from the multi-reference loss is considerably more regular.

4.5 Effect of the Time-Aware Flow

To assess the effect of the proposed time-aware warp (8), we conducted experiments on MVSEC, DSEC and ECD [28] datasets. Accuracy results are already reported in Tabs. 1 and 2. We now report values of the FWL metric in Tab. 3. For MVSEC, $dt = 1$ is a very short time interval, with small motion and therefore few events, hence the sharpness of the IWE with or without motion compensation are about the same (FWL ≈ 1). Instead, $dt = 4$ provides more events, and larger FWL values (1.1–1.3), which means that the contrast of the motion-compensated IWE is better than that of the zero flow baseline. All three meth-

Table 3. FWL (IWE sharpness) results on MVSEC, DSEC, and ECD. Higher is better.

	MVSEC ($dt = 4$)				ECD	DSEC	
	indoor1	indoor2	indoor3	outdoor1	slider_depth	thun_00_a	zurich_city_07_a
Ground truth	1.09	1.20	1.12	1.07	-	1.01	1.04
Ours: w/o time aware	1.17	1.30	1.23	1.11	1.88	1.39	1.57
Ours: Upwind	1.17	1.30	1.23	1.11	1.92	1.40	1.60
Ours: Burgers'	1.17	1.30	1.23	1.11	1.93	1.42	1.63

**Fig. 8.** *Time-aware flow.* Comparison between 3 versions of our method: Burgers', upwind, and no time-aware (4). At occlusions (dartboard in slider_depth [28] and garage door in DSEC [17]), upwind and Burgers' produce sharper IWEs. Due to the smoothness of the flow conferred by the tile-based approach, some small regions are still blurry.

ods provide sharper IWEs than ground truth. The advantages of the time-aware warp (8) over (4) to produce better IWEs (higher FWL) are most noticeable on sequences like slider_depth [28] and DSEC (see Fig. 8) because of the occlusions and larger motions. Notice that FWL differences below 0.1 are significant [45], demonstrating the efficacy of time-awareness.

4.6 Application to Deep Neural Networks (DNN)

The proposed secrets are not only applicable to model-based methods, but also to unsupervised-learning methods. We train EV-FlowNet [56] in an unsupervised manner, using (9) as data-fidelity term and a Charbonnier loss [8] as the regularizer. Since the time-aware flow does not have a significant influence on the MVSEC benchmark (Tab. 1), we do not port it to the learning-based setting. We convert 40k events into the voxel-grid representation [57] with 5 time bins. The network is trained for 50 epochs with a learning rate of 0.001 with Adam optimizer and with 0.8 learning rate decay. To ensure generalization, we train our network on indoor sequences and test on the outdoor_day1 sequence.

Table 4 shows the quantitative comparison with unsupervised-learning methods. Our model achieves the second best accuracy, following [57], and the best sharpness (FWL) among the existing methods. Notice that [57] was trained on the outdoor_day2 sequence, which is a similar driving sequence to the test one,

Table 4. Results of unsupervised learning on MVSEC’s outdoor_day1 sequence.

	$dt = 1$			$dt = 4$		
	AEE ↓	%Out ↓	FWL ↑	AEE ↓	%Out ↓	FWL ↑
EV-FlowNet [57]	0.32	0.00	-	1.30	9.70	-
EV-FlowNet (retrained) [33]	0.92	5.40	-	-	-	-
ConvGRU-EV-FlowNet [21]	0.47	0.25	0.94	1.69	12.50	0.94
Our EV-FlowNet using (9)	<u>0.36</u>	<u>0.09</u>	0.96	<u>1.49</u>	<u>11.72</u>	1.11

while the other methods were trained on drone data [9]. Hence [57] might be overfitting to the driving data, while ours is not, by the choice of training data.

4.7 Limitations

Like previous unsupervised works [21, 57], our method is based on the brightness constancy assumption. Hence, it struggles to estimate flow from events that are not due to motion, such as those caused by flickering lights. SL and SSL methods may forego this assumption, but they require high quality supervisory signal, which is challenging due to the HDR and high speed of event cameras.

Like other optical flow methods, our approach can suffer from the aperture problem. The flow could still collapse (events may be warped to too few pixels) if tiles become smaller (higher DOFs), or without proper regularization or initialization. Optical flow is also difficult to estimate in regions with few events, such as homogeneous brightness regions and regions with small apparent motion. Regularization fills in the homogeneous regions, whereas recurrent connections (like in RNNs) could help with small apparent motion.

5 Conclusion

We have extended the CM framework to estimate dense optical flow, proposing principled solutions to overcome problems of overfitting, occlusions and convergence without performing event voxelization. The comprehensive experiments show that our method achieves the best accuracy among all methods in the MVSEC indoor benchmark, and among the unsupervised and model-based methods in the outdoor sequence. It is also competitive in the DSEC optical flow benchmark. Moreover, our method delivers the sharpest IWEs and exposes the limitations of the benchmark data. Finally, we show how our method can be ported to the unsupervised setting, producing remarkable results. We hope our work unlocks future optical flow research on stable and interpretable methods.

Acknowledgements. We thank Prof. A. Yezzi and Dr. A. Zhu for useful discussions. Funded by the German Academic Exchange Service (DAAD), Research Grant - Bi-nationally Supervised Doctoral Degrees/Cotutelle, 2021/22 (57552338). Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2002/1 “Science of Intelligence” – project number 390523135.

References

1. Akolkar, H., Ieng, S.H., Benosman, R.: Real-time high speed motion prediction using fast aperture-robust event-driven visual flow. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**(1), 361–372 (2022). <https://doi.org/10.1109/TPAMI.2020.3010468>
2. Bardow, P., Davison, A.J., Leutenegger, S.: Simultaneous optical flow and intensity estimation from an event camera. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 884–892 (2016). <https://doi.org/10.1109/CVPR.2016.102>
3. Benosman, R., Clercq, C., Lagorce, X., Ieng, S.H., Bartolozzi, C.: Event-based visual flow. *IEEE Trans. Neural Netw. Learn. Syst.* **25**(2), 407–417 (2014). <https://doi.org/10.1109/TNNLS.2013.2273537>
4. Benosman, R., Ieng, S.H., Clercq, C., Bartolozzi, C., Srinivasan, M.: Asynchronous frameless event-based optical flow. *Neural Netw.* **27**, 32–37 (2012). <https://doi.org/10.1016/j.neunet.2011.11.001>
5. Brebion, V., Moreau, J., Davoine, F.: Real-time optical flow for vehicular perception with low- and high-resolution event cameras. *IEEE Trans. Intell. Transportation Systems* pp. 1–13 (2021). <https://doi.org/10.1109/TITS.2021.3136358>
6. Brosch, T., Tschechne, S., Neumann, H.: On event-based optical flow detection. *Front. Neurosci.* **9**(137) (Apr 2015). <https://doi.org/10.3389/fnins.2015.00137>
7. Cannici, M., Ciccone, M., Romanoni, A., Matteucci, M.: A differentiable recurrent surface for asynchronous event-based data. In: *Eur. Conf. Comput. Vis. (ECCV)*. pp. 136–152 (2020)
8. Charbonnier, P., Blanc-Feraud, L., Aubert, G., Barlaud, M.: Deterministic edge-preserving regularization in computed imaging. *IEEE Trans. Image Process.* **6**(2), 298–311 (1997). <https://doi.org/10.1109/83.551699>
9. Delmerico, J., Cieslewski, T., Rebecq, H., Faessler, M., Scaramuzza, D.: Are we ready for autonomous drone racing? the UZH-FPV drone racing dataset. In: *IEEE Int. Conf. Robot. Autom. (ICRA)*. pp. 6713–6719 (2019). <https://doi.org/10.1109/ICRA.2019.8793887>
10. Ding, Z., Zhao, R., Zhang, J., Gao, T., Xiong, R., Yu, Z., Huang, T.: Spatio-temporal recurrent networks for event-based optical flow estimation. *AAAI Conf. Artificial Intell.* **36**(1), 525–533 (2022)
11. Evans, L.C.: *Partial Differential Equations*. Graduate Studies in Mathematics, American Mathematical Society (2010)
12. Gallego, G., Delbruck, T., Orchard, G., Bartolozzi, C., Tabá, B., Censi, A., Leutenegger, S., Davison, A., Conradt, J., Daniilidis, K., Scaramuzza, D.: Event-based vision: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**(1), 154–180 (2022). <https://doi.org/10.1109/TPAMI.2020.3008413>
13. Gallego, G., Gehrig, M., Scaramuzza, D.: Focus is all you need: Loss functions for event-based vision. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 12272–12281 (2019). <https://doi.org/10.1109/CVPR.2019.01256>
14. Gallego, G., Rebecq, H., Scaramuzza, D.: A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 3867–3876 (2018). <https://doi.org/10.1109/CVPR.2018.00407>
15. Gallego, G., Scaramuzza, D.: Accurate angular velocity estimation with an event camera. *IEEE Robot. Autom. Lett.* **2**(2), 632–639 (2017). <https://doi.org/10.1109/LRA.2016.2647639>
16. Gehrig, D., Loquercio, A., Derpanis, K.G., Scaramuzza, D.: End-to-end learning of representations for asynchronous event-based data. In: *Int. Conf. Comput. Vis. (ICCV)*. pp. 5632–5642 (2019). <https://doi.org/10.1109/ICCV.2019.00573>

17. Gehrig, M., Aarents, W., Gehrig, D., Scaramuzza, D.: DSEC: A stereo event camera dataset for driving scenarios. *IEEE Robot. Autom. Lett.* **6**(3), 4947–4954 (2021). <https://doi.org/10.1109/LRA.2021.3068942>
18. Gehrig, M., Millhäusler, M., Gehrig, D., Scaramuzza, D.: E-RAFT: Dense optical flow from event cameras. In: *Int. Conf. 3D Vision (3DV)*. pp. 197–206 (2021). <https://doi.org/10.1109/3DV53792.2021.00030>
19. Gonzalez, R.C., Woods, R.E.: *Digital Image Processing*. Pearson Education (2009)
20. Gu, C., Learned-Miller, E., Sheldon, D., Gallego, G., Bideau, P.: The spatio-temporal Poisson point process: A simple model for the alignment of event camera data. In: *Int. Conf. Comput. Vis. (ICCV)*. pp. 13495–13504 (2021). <https://doi.org/10.1109/ICCV48922.2021.01324>
21. Hagenaaars, J.J., Paredes-Valles, F., de Croon, G.C.H.E.: Self-supervised learning of event-based optical flow with spiking neural networks. In: *Advances in Neural Information Processing Systems (NeurIPS)*. vol. 34, pp. 7167–7179 (2021). <https://doi.org/10.48550/arXiv.2106.01862>
22. Kim, H., Kim, H.J.: Real-time rotational motion estimation with contrast maximization over globally aligned events. *IEEE Robot. Autom. Lett.* **6**(3), 6016–6023 (2021). <https://doi.org/10.1109/LRA.2021.3088793>
23. Lagorce, X., Orchard, G., Gallupi, F., Shi, B.E., Benosman, R.: HOTS: A hierarchy of event-based time-surfaces for pattern recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(7), 1346–1359 (Jul 2017). <https://doi.org/10.1109/TPAMI.2016.2574707>
24. Lee, C., Kosta, A., Zhu, A.Z., Chaney, K., Daniilidis, K., Roy, K.: Spike-flownet: Event-based optical flow estimation with energy-efficient hybrid neural networks. In: *Eur. Conf. Comput. Vis. (ECCV)*. pp. 366–382 (2020)
25. Li, H., Li, G., Shi, L.: Super-resolution of spatiotemporal event-stream image. *Neurocomputing* **335**, 206–214 (2019). <https://doi.org/10.1016/j.neucom.2018.12.048>
26. Liu, M., Delbruck, T.: Adaptive time-slice block-matching optical flow algorithm for dynamic vision sensors. In: *British Mach. Vis. Conf. (BMVC)*. pp. 1–12 (2018)
27. Mitrokhin, A., Fermuller, C., Parameshwara, C., Aloimonos, Y.: Event-based moving object detection and tracking. In: *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*. pp. 1–9 (2018). <https://doi.org/10.1109/IROS.2018.8593805>
28. Mueggler, E., Rebecq, H., Gallego, G., Delbruck, T., Scaramuzza, D.: The event-camera dataset and simulator: Event-based data for pose estimation, visual odometry, and SLAM. *Int. J. Robot. Research* **36**(2), 142–149 (2017). <https://doi.org/10.1177/0278364917691115>
29. Nagata, J., Sekikawa, Y., Aoki, Y.: Optical flow estimation by matching time surface with event-based cameras. *Sensors* **21**(4) (2021). <https://doi.org/10.3390/s21041150>
30. Nunes, U.M., Demiris, Y.: Robust event-based vision model estimation by dispersion minimisation. *IEEE Trans. Pattern Anal. Mach. Intell.* (2021). <https://doi.org/10.1109/TPAMI.2021.3130049>
31. Orchard, G., Benosman, R., Etienne-Cummings, R., Thakor, N.V.: A spiking neural network architecture for visual motion estimation. In: *IEEE Biomed. Circuits Syst. Conf. (BioCAS)*. pp. 298–301 (2013). <https://doi.org/10.1109/biocas.2013.6679698>
32. Parameshwara, C.M., Sanket, N.J., Singh, C.D., Fermüller, C., Aloimonos, Y.: 0-MMS: Zero-shot multi-motion segmentation with a monocular event camera. In: *IEEE Int. Conf. Robot. Autom. (ICRA)*. pp. 9594–9600 (2021). <https://doi.org/10.1109/ICRA48506.2021.9561755>

33. Paredes-Valles, F., de Croon, G.C.H.E.: Back to event basics: Self-supervised learning of image reconstruction for event cameras via photometric constancy. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 3445–3454 (2021). <https://doi.org/10.1109/CVPR46437.2021.00345>
34. Peng, X., Gao, L., Wang, Y., Kneip, L.: Globally-optimal contrast maximisation for event cameras. *IEEE Trans. Pattern Anal. Mach. Intell.* **44**(7), 3479–3495 (2022). <https://doi.org/10.1109/TPAMI.2021.3053243>
35. Posch, C., Serrano-Gotarredona, T., Linares-Barranco, B., Delbruck, T.: Retinomorph event-based vision sensors: Bioinspired cameras with spiking output. *Proc. IEEE* **102**(10), 1470–1484 (Oct 2014). <https://doi.org/10.1109/jproc.2014.2346153>
36. 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 (Dec 2018). <https://doi.org/10.1007/s11263-017-1050-6>
37. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 234–241 (2015)
38. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* **60**(1–4), 259–268 (Nov 1992). [https://doi.org/10.1016/0167-2789\(92\)90242-F](https://doi.org/10.1016/0167-2789(92)90242-F)
39. Seok, H., Lim, J.: Robust feature tracking in dvs event stream using Bezier mapping. In: *IEEE Winter Conf. Appl. Comput. Vis. (WACV)*. pp. 1647–1656 (2020). <https://doi.org/10.1109/WACV45572.2020.9093607>
40. Sethian, J.: *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*. Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press (1999)
41. Shiba, S., Aoki, Y., Gallego, G.: Event collapse in contrast maximization frameworks. *Sensors* **22**(14), 1–20 (2022). <https://doi.org/10.3390/s22145190>
42. Stoffregen, T., Gallego, G., Drummond, T., Kleeman, L., Scaramuzza, D.: Event-based motion segmentation by motion compensation. In: *Int. Conf. Comput. Vis. (ICCV)*. pp. 7243–7252 (2019). <https://doi.org/10.1109/ICCV.2019.00734>
43. Stoffregen, T., Kleeman, L.: Simultaneous optical flow and segmentation (SOFAS) using Dynamic Vision Sensor. In: *Australasian Conf. Robot. Autom. (ACRA)* (2017)
44. Stoffregen, T., Kleeman, L.: Event cameras, contrast maximization and reward functions: an analysis. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 12292–12300 (2019). <https://doi.org/10.1109/CVPR.2019.01258>
45. Stoffregen, T., Scheerlinck, C., Scaramuzza, D., Drummond, T., Barnes, N., Kleeman, L., Mahony, R.: Reducing the sim-to-real gap for event cameras. In: *Eur. Conf. Comput. Vis. (ECCV)*. pp. 534–549 (2020). https://doi.org/https://doi.org/10.1007/978-3-030-58583-9_32
46. Sun, D., Roth, S., Black, M.J.: A quantitative analysis of current practices in optical flow estimation and the principles behind them. *Int. J. Comput. Vis.* **106**(2), 115–137 (Sep 2013). <https://doi.org/10.1007/s11263-013-0644-x>
47. Taverni, G., Moeys, D.P., Li, C., Cavaco, C., Motsnyi, V., Bello, D.S.S., Delbruck, T.: Front and back illuminated Dynamic and Active Pixel Vision Sensors comparison. *IEEE Trans. Circuits Syst. II* **65**(5), 677–681 (2018). <https://doi.org/10.1109/TCSII.2018.2824899>

48. Teed, Z., Deng, J.: RAFT: Recurrent all pairs field transforms for optical flow. In: *Eur. Conf. Comput. Vis. (ECCV)*. pp. 402–419 (2020). https://doi.org/10.1007/978-3-030-58536-5_24
49. Trucco, E., Verri, A.: *Introductory Techniques for 3-D Computer Vision*. Prentice Hall PTR, Upper Saddle River, NJ, USA (1998)
50. Ye, C., Mitrokhin, A., Parameshwara, C., Fermüller, C., Yorke, J.A., Aloimonos, Y.: Unsupervised learning of dense optical flow, depth and egomotion with event-based sensors. In: *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*. pp. 5831–5838 (2020). <https://doi.org/10.1109/IROS45743.2020.9341224>
51. Zhang, Z., Yezzi, A., Gallego, G.: Image reconstruction from events. Why learn it? *arXiv e-prints* (2021). <https://doi.org/10.48550/arXiv.2112.06242>
52. Zhou, Y., Gallego, G., Lu, X., Liu, S., Shen, S.: Event-based motion segmentation with spatio-temporal graph cuts. *IEEE Trans. Neural Netw. Learn. Syst.* pp. 1–13 (2021). <https://doi.org/10.1109/TNNLS.2021.3124580>
53. Zhu, A.Z., Atanasov, N., Daniilidis, K.: Event-based feature tracking with probabilistic data association. In: *IEEE Int. Conf. Robot. Autom. (ICRA)*. pp. 4465–4470 (2017). <https://doi.org/10.1109/ICRA.2017.7989517>
54. Zhu, A.Z., Atanasov, N., Daniilidis, K.: Event-based visual inertial odometry. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 5816–5824 (2017). <https://doi.org/10.1109/CVPR.2017.616>
55. Zhu, A.Z., Thakur, D., Ozaslan, T., Pfrommer, B., Kumar, V., Daniilidis, K.: The multivehicle stereo event camera dataset: An event camera dataset for 3D perception. *IEEE Robot. Autom. Lett.* **3**(3), 2032–2039 (Jul 2018). <https://doi.org/10.1109/lra.2018.2800793>
56. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: EV-FlowNet: Self-supervised optical flow estimation for event-based cameras. In: *Robotics: Science and Systems (RSS)* (2018). <https://doi.org/10.15607/RSS.2018.XIV.062>
57. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: Unsupervised event-based learning of optical flow, depth, and egomotion. In: *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. pp. 989–997 (2019). <https://doi.org/10.1109/CVPR.2019.00108>