

# EVReflex: Dense Time-to-Impact Prediction for Event-based Obstacle Avoidance

Celyn Walters<sup>1</sup> and Simon Hadfield<sup>1</sup>

**Abstract**— The broad scope of obstacle avoidance has led to many kinds of computer vision-based approaches. Despite its popularity, it is not a solved problem. Traditional computer vision techniques using cameras and depth sensors often focus on static scenes, or rely on priors about the obstacles. Recent developments in bio-inspired sensors present event cameras as a compelling choice for dynamic scenes. Although these sensors have many advantages over their frame-based counterparts, such as high dynamic range and temporal resolution, event-based perception has largely remained in 2D. This often leads to solutions reliant on heuristics and specific to a particular task.

We show that the fusion of events and depth overcomes the failure cases of each individual modality when performing obstacle avoidance. Our proposed approach unifies event camera and lidar streams to estimate metric Time-To-Impact (TTI) without prior knowledge of the scene geometry or obstacles. In addition, we release an extensive event-based dataset with six visual streams spanning over 700 scanned scenes.

## I. INTRODUCTION

Obstacle avoidance is a cornerstone of autonomous robotic behaviour. It enables robots to explore unknown and dynamic environments safely, as well as being vital for Human-Computer Interaction (HCI). As such, it has been a topic of interest in the field of computer vision since its inception. Vision-based techniques have historically relied on salient cues such as image discontinuities and corner features. This causes vision-based approaches to perform less well in poor lighting or where there is little or repetitive texture. In these scenarios, many motion planning implementations utilise depth sensors like lidar, which mitigate issues with texture, lighting and motion blur. However, limited resolution and slow scan times make it challenging to evade small or fast moving objects. This is particularly an issue in domains like autonomous driving, where fast reaction times are vital.

Event cameras are an emerging technology which are a promising alternative or supplement to traditional frame-based cameras and depth sensors. Rather than producing entire image frames, event cameras output a stream of per-pixel intensity changes. This has much higher temporal resolution and dynamic range than regular RGB images, and almost eliminates motion blur. The bandwidth of the event stream can widen or narrow depending on the amount of change observed. These properties are clearly useful for the task of

\* This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant agreement EP/S035761/1 "Reflexive robotics". We would also like to thank Dr. Hamideh Kerdegari for her help and initial work on this topic.

<sup>1</sup>Celyn Walters and Simon Hadfield are with CVSSP, University of Surrey, UK celyn.walters@surrey.ac.uk




Fig. 1: Illustration of sensor capabilities for collision avoidance. Depth cameras work well for large static objects but struggle to detect small fast-moving objects. Event cameras can segment dynamic foreground objects but cannot infer depth to avoid large textureless obstacles. Our work demonstrates their fusion works in both cases.

obstacle avoidance, but due to the different data composition, specialised techniques must be created to utilise these cameras.

Similar to traditional frame-based cameras, event cameras operate on a 2D image plane and require dedicated structure from motion [1] or Machine Learning (ML) techniques [2, 3] to infer depth, introducing delays. The lack of direct depth estimation is a shortcoming for obstacle avoidance, as it is often not possible to calculate the time or location of potential collisions, and as a result, the danger to be associated to given objects. Additionally, as with frame based cameras, it is impossible for an event camera to observe flat coloured textureless objects like walls. The lack of intensity changes causes no events to be generated in these regions.

Similar to human stereo vision, the complementary properties of events and depth have the potential to enable robust obstacle avoidance across a previously infeasible range of scenarios. Much like a biological creature, we argue that it is also important for the agent to be able to prioritize either depth perception or intensity changes depending on the situation.

As exemplified by Figure 1, this enables an agent to avoid obstacles with low texture and also those moving quickly, which would be challenging for a single sensor modality. To this end, we make the following contributions:

- 1) An extensive dataset for developing event-based computer vision techniques<sup>1</sup>: This includes RGB images, events, as well as ground truth depth, optical flow, semantic masks, and dense inverse TTI maps. The extracted data covers over 700 scanned scenes from the ScanNet dataset [4] with the inclusion of additional flying obstacles. The data are formatted both as image databases and as ROS messages.
- 2) To our knowledge, the first approach to unify lidar and event camera sensor streams within a deep-learning framework, as shown in Figure 1,
- 3) Demonstration that the above approach can be used to develop systems which avoid a broader range of dangerous obstacles than any previous technique.

## II. RELATED WORK

**Obstacle avoidance** — Fundamentally, solving obstacle avoidance in static scenes comprises of pose estimation and mapping. As such, robotic obstacle avoidance has shared roots with visual odometry and Simultaneous Localization And Mapping (SLAM). The differences between them are often only in the framing of the problem. Early approaches such as the work of Borenstein and Koren utilise a certainty grid and potential field for collision-free local path planning [5]. More recently, approaches such as the work of Chakravarty et al. use depth estimation from single images [6]. These perform reasonably well in clear open environments but tend to over-smooth small and distinct details. The majority of approaches only consider static scenes, in which RGB and depth sensors are suitable. We propose that the fusion of depth with an event stream is more suitable where there is a mixture of static and dynamic obstacles.

**Events** — While the event representation has inherent advantages, it must be processed differently from traditional images. Event-based approaches for many classic computer vision problems have emerged recently. Zhu et al. estimate optical flow from events, using a self-supervised network supervised by corresponding greyscale images [7]. In Hidalgo-Carrió et al.’s work, a monocular event stream is used in a recurrent network to estimate depth using temporal consistency. EVDodge from Sanket et al. is an approach which uses events to estimate ‘segmentation flow’, optical flow with foreground/background segmentation [8]. They train in simulation and show a drone capable of dodging thrown objects.

**Event camera datasets** — Following the recent availability of event cameras, a few datasets have emerged each with a slightly different focus. For visual-inertial odometry (VIO), Mueggler et al. introduced an event-based dataset which includes images and IMU, with ground truth poses [9]. The

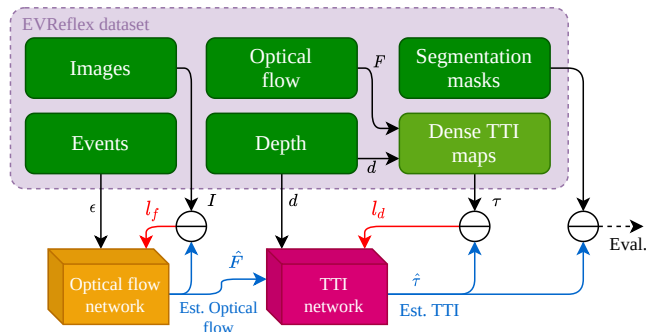


Fig. 2: Dataset and network data flow. The  $\ominus$  represents the application of the relevant loss function. Loss signals are notated in red, output signals are notated in blue.

popular Multivehicle Stereo Event Camera (MVSEC) dataset from Zhu et al. followed, which involves stereo event cameras and additionally provides lidar point clouds and more scene variety [10]. These have enabled successful approaches [7, 3, 2], but they lack extra data streams like ground truth optical flow and semantic labels.

Simulated datasets, while having a domain gap to be overcome for real-world applications, can provide accurate ground truth to supervise or evaluate other approaches. Early simulators emulated events by processing differences between image frames [11]. Mueggler et al. used a custom rendering engine for accurate event data in simple scenes [9]. InteriorNet from Li et al. boasts many additional streams such as optical flow and surface normals, in realistic customisable scenes [12]. However, [9, 12] still operate on frames. Rebecq et al. enhanced [9] with ESIM, which uses adaptive-rate sampling [13] but does not match the realism and flexibility of InteriorNet. We use real scanned scenes from the ScanNet dataset [4] in conjunction with ESIM for a varied dataset which correctly models an event camera.

## III. METHODOLOGY

As shown in Figure 2, our approach consists of two parts. Firstly, the event stream  $\epsilon$  is used to predict optical flow  $\hat{F}$  between image frames. The optical flow network  $O$  is self-supervised using the greyscale image stream  $I$  which corresponds with the events. In the second part, the resulting optical flow is combined with depth  $d$  and processed by the second subnetwork  $P$ . The TTI estimation network directly estimates inverse metric TTI, the supervisory signal being formed from depth frames warped by optical flow.

### A. Optical flow estimation

The task of optical flow estimation has had a long history of approaches based on traditional visible images. Inherent drawbacks with the characteristics of this data, such as motion blur, fixed framerates and poor low light performance, affect all of these approaches. While many hand-crafted techniques mitigate these problems [14, 15], event cameras naturally perceive small changes with high temporal resolution. The performance of event cameras for optical flow estimation has been shown to exceed traditional cameras particularly for thin

<sup>1</sup> The dataset, code, and instructions for running and generating additional data are available at <https://gitlab.eps.surrey.ac.uk/cw0071/EVReflex>




Fig. 3: Flow estimation network architecture. The  $\otimes$  represents a concatenation operation. Total loss is calculated as the mean of losses from each decoder layer’s output  $F_{1-4}$ .

structures and low-texture areas [7]. These advantages are particularly important for obstacle avoidance, where we can expect fast-moving objects to exhibit motion blur in traditional images.

Inspired by the work of Zhu et al. [7], the first module of our approach is a network which estimates the optical flow for use in our second module (Section III-B). This network is trained in a self-supervised manner, using the grayscale images in the training loss function. This means that ground truth optical flow is not required to train for a new dataset.

Although events themselves are streamed asynchronously between pixels, many works opt to treat them similar to image frames. This is often implemented as an accumulation of events between  $t$  and  $t - 1$ . It is known that both the positive/negative event counts and event times are informative signals for event-based learning [16]. Following the work of Zhu et al., we represent the accumulated events  $\epsilon$  as a 4 channel image  $[4, H, W]$ . The channels of this event map are the positive and negative event count at that pixel ( $\epsilon^{n+}$ ,  $\epsilon^{n-}$ ) and the timestamps of the most recent positive/negative event ( $\epsilon^{t+}$ ,  $\epsilon^{t-}$ ). This compact format has been shown to be sufficient for many event-based image processing tasks [17, 18]. Resembling a U-Net [19], our network  $O$  is composed of four encoder and four decoder layers with skip connections, shown in Figure 3. Given network weights  $\theta_O$ , the result is an estimated flow field  $F$  at each decoder layer:

$$F = O(\epsilon|\theta_O). \quad (1)$$

The application of  $F$  on the corresponding grayscale image  $I_t$  is compared with the next image  $I_{t+1}$  using photometric and smoothness losses. The photometric loss  $l_p$  minimises intensity differences with Charbonnier loss [20]  $\rho$ :

$$l_p(F, I_t, I_{t+1}) = \sum_i \rho(I_t(i) - I_{t+1}(\mathcal{W}(i|F))) \quad \forall i \in I, \quad (2)$$

where  $i$  is a pixel coordinate  $\{x, y\}$  and  $\mathcal{W}$  warps an image by sampling it according to a flow field.  $\rho$  is calculated using  $\epsilon=0.001, \alpha=0.45$ . The smoothness loss  $l_s$  regularises flow discontinuities in a pixel’s neighbourhood  $N$ ,

$$l_s(F) = \sum_i \sum_{n \in N} \rho(i - n) \quad \forall i \in F. \quad (3)$$




Fig. 4: Inverse TTI estimation network architecture.

The two losses are combined as a weighted sum,  $l_f = l_p + \alpha l_s$ , with  $\alpha$  as the weight. By substituting Equations (2) and (3), the training process for the self-supervised flow network in isolation is

$$\hat{\theta}_O = \operatorname{argmin}_{\theta_O} (l_p(O(\epsilon|\theta_P), I_t, I_{t+1}) + \alpha l_s(O(\epsilon|\theta_O))). \quad (4)$$

### B. TTI estimation

The task of this network is to regress the inverse TTI  $\hat{\tau}$  for each pixel given optical flow  $F$  and depth values  $d$ . Inverse TTI is used to avoid the asymptote in TTI at the point where the relative egocentric velocity between the agent and an obstacle is zero. To investigate the network’s reliance between optical flow and instantaneous depth, we define two versions, and each are tested independently in Section IV. The first, which we refer to as the ‘static’ version, takes the stacked optical flow  $F_t$  and instantaneous depth  $d_t$  as input. Here, the speed of geometry moving into the image plane must be estimated implicitly during the computation of the TTI map. The second version additionally uses the next sequential depth frame  $d_{t+1}$  to provide additional cues about motion in the Z axis. We refer to this version as the ‘dynamic’ version.

The network  $P$ , shown in Figure 4, downsamples the input twice through 9 convolutional layers with branched residual connections, and performs the transposed version to upsample back to the original resolution. The network is supervised by the ground truth inverse TTI  $\tau$ , explained in Section III-D, using mean-squared error loss  $l_\tau$  such that

$$\hat{\theta}_P = \operatorname{argmin}_{\theta_P} (l_\tau(\hat{\tau}, \tau)) \quad \text{where } \hat{\tau} = P([F, d]|\theta_P). \quad (5)$$

### C. Control policy

In the work of Sanket et al., some obstacle detection is based on priors. For completely unknown objects, the input to the control policy is the optical flow after masking segmented background regions [8]. The obstacle trajectory is assumed to be the average direction of the masked optical flow on the image plane, and evasion is directed to be perpendicular to this. There is no check that this motion is safe; the priority is to avoid a fast-moving object.

For the comparison in Section IV-C, we take a similar approach without any priors on the obstacles. We combine the depth with predicted TTI into a 3D motion vector, and then use the cross-product of this with the agent’s egomotion

$v$  as the direction of evasion  $\psi$ ,

$$\psi = \left( \frac{1}{n} \sum_{i=1}^n (F(i), d(i)\hat{\tau}(i)) \right) \times v. \quad (6)$$

Because our network makes a prediction for every pixel, we can additionally avoid colliding with ‘background’ regions if they are within the field of view.

#### D. Dataset

A dataset, EVReflex, was created to facilitate evaluation of our approach. With its extensive ground truth for multiple sensors, it is also a useful asset for future event camera-based research.

To collect this dataset, the movement of a robotic agent is simulated in indoor scenes. The robot is equipped with an RGB camera, a depth camera, IMU, and an event camera. The simulator also outputs the pose of the agent and ground truth optical flow [13], and has been extended to be capable of producing semantic class labels. At the time of writing, there are very few other datasets which provide an event stream as well as semantic labels.

A problem for many synthetic datasets is the domain gap between simulated environments and the real world, due to the lack of realism. This issue applies to most event camera datasets including EVDodge [8] and those based on ESIM [13]. EVReflex overcomes this major problem by simulating in RGB-D scanned scenes from the ScanNet dataset [4], with a total of over 700 different real-world rooms. The simulated robot follows random floor-based trajectories<sup>2</sup> within each room, which brings it into proximity with various static obstacles. Concurrent with this, a few dynamic obstacles are simulated, each of which follow randomized trajectories that periodically bring them into proximity with the robot.

A notable contribution of this dataset to aid future research on collision avoidance tasks is the additional image stream for the identification of dangerous image regions. Each pixel  $i$  in these images represents inverse Time-To-Impact (TTI) values for that point in the scene. These are computed by comparing the previous depth  $d_{t-1}$  warped by the optical flow  $F$  with the current depth  $d_t$ ,

$$\tau(i) = \max \left( 0, \frac{d_t(i) - d_{t-1}(\mathcal{W}(i|F))}{d_t} \right). \quad (7)$$

This is fundamentally different from foreground/background segmentation such as the estimated in EVDodge [8]. Firstly, this approach can function with an arbitrary number of obstacles of any shape. There is no intrinsic preference between compact dynamic obstacles and static geometry. Secondly, objects which are not approaching the robot, including those in close proximity, have zero values.

The dataset is available to download<sup>1</sup> in both image databases and ROS formats. We also provide the full simulation system to allow people to customise and generate additional data if needed.

<sup>2</sup>Most translational variation is in the  $[X, Y]$  floor plane, and most rotational variation in  $Z$  rotation (yaw).

## IV. EXPERIMENTS

The evaluation is split into three parts. Firstly, evaluation was performed on the performance of the optical flow network. Secondly the accuracy of the dense TTI regression. This evaluation was broken down across different semantic classes to highlight the performance for different types of obstacles. Where applicable, we compare against the state-of-the-art in event-based obstacle avoidance (EVDodge [8]), and against a depth-sensor baseline, which thresholds nearby objects as obstacles. Finally, we evaluate the accuracy of the signal for the control policy across the baselines.

The EVReflex dataset was split into training, evaluation and test splits at a ratio of 70%, 15% and 15% respectively. This split was performed across the environments, meaning that the same environment is not visible in more than one set.

The optical flow estimation subnetwork (Section III-A) was trained with a initial learning rate of  $1 \times 10^{-5}$ , with a batch size of 8 for 30 epochs.

The TTI estimation subnetwork (Section III-B) was trained for 40 epochs, with a batch size of 40 and a learning rate of 0.01 which reduced by a factor of 10 on a validation loss plateau.

#### A. Event-based optical flow evaluation

We evaluate optical flow by computing average endpoint error (EE) with respect to the ground truth optical flow. By filtering to event locations only, the system achieved an overall average EE of 0.36, with 0.47% of those being outliers with an EE of greater than 3 pixels. However, without filtering, the average EE is 15.69 with 84.44% outliers. Our dataset utilises laser-scanned scenes from the ScanNet dataset. The scanning process has high spatial accuracy, but textures are smoothed. This leads to decreased events, and as a result, higher difficulty estimating optical flow. This is especially visible in Figure 5. Note that regions with events are still generally correct.

This reliance on salient texture is a drawback of visual sensors in general, and is one of the motivations of our overall approach. Depth sensors perceive textureless objects well, and Section IV-B shows that fusing instantaneous depth with this optical flow is enough to determine potential collisions.

#### B. Dense TTI evaluation

A naïve approach to obstacle avoidance may be to consider the closest objects as the most dangerous. This may be true for linear forward motion, but does not necessarily hold true for others. We show a qualitative comparison of the estimated TTI versus the ground truth and EVDodge (SegFlowNet) [8] in Figure 6. EVDodge expects a clear foreground/background distinction and without it the performance is reduced. Where the camera is rotating clockwise, static geometry on the left side of the image draws nearer while that on the right side moves further away. Similarly, where the camera is moving backwards, TTI does not correlate with depth. For these cases, instantaneous depth is insufficient for collision avoidance. The output of our approach resembles the ground truth TTI rather




Fig. 5: Qualitative optical flow estimation results. Clearly, optical flow cannot be estimated for low texture regions which do not generate events.




Fig. 6: Qualitative TTI estimation results. **Row 1:** Camera is moving forwards, ball is approaching camera. **Row 2:** Camera is rotating clockwise. **Row 3:** Camera is moving backwards, TTI is negative.

than depth, and it has learned to associate motion derived from events with the depth data. Note that we do not use explicit pose information.

To enable quantitative comparison with previous techniques based on binary segmentation of dangerous objects, we threshold our inverse TTI values to pixels which are projected to collide within 1.0s. We compare both versions of our network against the naïve thresholded depth baseline using a value of 0.5m. This value was chosen to give similar output to the other approaches. We also compare against EVDodge, current state-of-the-art for event-based obstacle avoidance [8], for which we similarly apply a binary threshold to its ‘segmentation flow’ output. Table I shows that our approach more consistently identifies the dangerous regions. The dataset, and as such, the task, is not suited to EVDodge, which expects clearer foreground/background delineation. The ‘dynamic’ version of our approach achieves very high precision

but poor recall, likely as an over-reliance on depth cues. The regular ‘static’ version outperforms the baselines by a wide margin for recall and  $f_1$  scores. The performance of all approaches is better for flying objects, which produce many events, than on the floor class, which produces few events.

Approach	Flying objects			Floor			Overall		
	$p$	$r$	$f_1$	$p$	$r$	$f_1$	$p$	$r$	$f_1$
Depth baseline	0.88	0.82	0.85	0.69	0.26	0.38	0.55	0.27	0.36
EVDodge	0.35	0.40	0.37	0.87	0.27	0.42	0.75	0.22	0.34
Ours	0.94	<b>0.89</b>	<b>0.91</b>	0.38	<b>0.95</b>	<b>0.54</b>	0.61	<b>0.95</b>	<b>0.74</b>
Ours (dynamic)	<b>1.00</b>	0.59	0.74	<b>0.99</b>	0.25	0.41	<b>0.99</b>	0.22	0.36

TABLE I: Precision ( $p$ ), recall ( $r$ ) and F1 ( $f_1$ ) scores for thresholded TTI (potential collision within 1 second), with semantic-specific scores.

### C. Evasion

Where Section IV-B examined the magnitude and spatial accuracy of the predicted inverse TTI maps, a control policy such as described in Section III-C acts upon directional cues. The evasion direction is based on the estimated motion vector of obstacles relative to the agent. With a fixed control policy, the accuracy of this estimate represents an approach’s ability to avoid oncoming obstacles. Over the test set, we evaluate the average error magnitude between the estimated and ground truth motion vectors, and report the values in Table II. The table shows that all our approach much more

Approach	AAE	AAE (top 10%)
Inverse depth (baseline)	16.58°	27.54°
EVDodge	25.80°	22.28°
Ours	<b>2.65°</b>	<b>4.93°</b>
Ours (dynamic)	10.70°	17.11°

TABLE II: Average angle error for estimated trajectory of dangerous obstacles. ‘Top 10%’ represents the values with the top 10% motion vector magnitudes

accurately estimates the direction with the most associated danger. Those in the top 10% of ground truth magnitudes are particularly important as this relates to having the least time to avoid a collision. This accuracy is possible through being able to filter out objects which may be in close proximity but are not currently dangerous.

## V. CONCLUSIONS

In this paper, we proposed that the fusion of events and depth data enable better collision avoidance strategies than using a single modality. Depth sensors, while accurate for simple static scenes, struggle with fast-moving objects. On the other hand, event cameras have excellent temporal resolution and dynamic range, but are not effective with textureless surfaces. Each sensor modality complements the failure cases of the other. We presented an approach which overcomes these issues, and directly estimates Time-To-Impact (TTI) for every pixel. Our method outperforms those based on foreground/background segmentation, and we demonstrated

that the resulting TTI maps can subsequently be used to direct a control policy to circumvent collisions. We believe that this is an important step in utilising the benefits of bio-inspired sensors. To complement our approach, we provide an extensive supporting dataset with a wide variety of data types to support further research.

#### REFERENCES

- [1] Hanme Kim, Stefan Leutenegger and Andrew J. Davison. ‘Real-time 3D reconstruction and 6-DoF tracking with an event camera’. In: *Lect. Notes Comput. Sci. (LNCS)* 9910 (2016), pp. 349–364.
- [2] Javier Hidalgo-Carrió, Daniel Gehrig and Davide Scaramuzza. ‘Learning Monocular Dense Depth from Events’. In: *IEEE Int. Conf. 3D Vision (3DV)*. 2020.
- [3] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney and Kostas Daniilidis. ‘Unsupervised event-based learning of optical flow, depth, and egomotion’. In: *Proc. IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*. 2019, pp. 989–997.
- [4] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser and Matthias Nießner. ‘ScanNet: Richly-annotated 3D reconstructions of indoor scenes’. In: *Proc. IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*. 2017, pp. 2432–2443.
- [5] Johann Borenstein and Yorem Koren. ‘Real-Time Obstacle Avoidance for Fast Mobile Robots’. In: *IEEE Trans. Syst. Man Cybern. (TSMC)* 19.5 (1989), pp. 1179–1187.
- [6] Punarjay Chakravarty, Klaas Kelchtermans, Tom Rousel, Stijn Wellens, Tinne Tuytelaars and Luc Van Eycken. ‘CNN-based single image obstacle avoidance on a quadrotor’. In: *Proc. IEEE Int. Conf. Robot. and Autom. (ICRA)*. IEEE, 2017, pp. 6369–6374.
- [7] Alex Zhu, Liangzhe Yuan, Kenneth Chaney and Kostas Daniilidis. ‘EV-FlowNet: Self-Supervised Optical Flow Estimation for Event-based Cameras’. In: *Robot. Sci. Syst. XIV. Robotics: Science and Systems Foundation*, June 2018.
- [8] Nitin J. Sanket, Chethan M. Parameshwara, Chahat Deep Singh, Ashwin V. Kuruttukulam, Cornelia Fermüller, Davide Scaramuzza and Yiannis Aloimonos. ‘EVDodge: Embodied AI For High-Speed Dodging On A Quadrotor Using Event Cameras’. In: *Proc. IEEE Int. Conf. Robot. and Autom. (ICRA)*. 2019.
- [9] Elias Mueggler, Henri Rebecq, Guillermo Gallego, Tobi Delbruck and Davide Scaramuzza. ‘The event-camera dataset and simulator: Event-based data for pose estimation, visual odometry, and SLAM’. In: *Int. J. Rob. Res. (IJRR)* 36.2 (2017), pp. 142–149.
- [10] Alex Zihao Zhu, Dinesh Thakur, Tolga Ozaslan, Bernd Pfrommer, Vijay Kumar and Kostas Daniilidis. ‘The Multivehicle Stereo Event Camera Dataset: An Event Camera Dataset for 3D Perception’. In: *IEEE Robot. Autom. Lett. (RA-L)* 3.3 (July 2018), pp. 2032–2039.
- [11] Jacques Kaiser, J. C.V. Tieck, Christian Hubschneider, Peter Wolf, Michael Weber, Michael Hoff, Alexander Friedrich, Konrad Wojtasik, Arne Roennau, Ralf Kohlhaas, Rudiger Dillmann and J. Marius Zollner. ‘Towards a framework for end-to-end control of a simulated vehicle with spiking neural networks’. In: *Proc. IEEE Int. Conf. Simulation, Model. Program. Auton. Robot. (SIMPAN)* (2017), pp. 127–134.
- [12] Wenbin Li, Sajad Saeedi, John McCormac, Ronald Clark, Dimos Tzoumanikas, Qing Ye, Yuzhong Huang, Rui Tang and Stefan Leutenegger. ‘InteriorNet: Mega-scale multi-sensor photo-realistic indoor scenes dataset’. In: *British Machine Vision Conference (BMVC)*. 2018.
- [13] Henri Rebecq, Daniel Gehrig and Davide Scaramuzza. ‘ESIM: an Open Event Camera Simulator’. In: *Proc. 2nd Conf. Robot Learn.* 2018, pp. 969–982.
- [14] Philippe Weinzaepfel, Jerome Revaud, Zaid Harchaoui and Cordelia Schmid. ‘DeepFlow: Large Displacement Optical Flow with Deep Matching’. In: *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*. IEEE, Dec. 2013, pp. 1385–1392.
- [15] Jerome Revaud, Philippe Weinzaepfel, Zaid Harchaoui and Cordelia Schmid. ‘EpicFlow: Edge-preserving interpolation of correspondences for optical flow’. In: *Proc. IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*. 2015, pp. 1164–1172.
- [16] Anh Nguyen, Thanh Toan Do, Darwin G. Caldwell and Nikos G. Tsagarakis. ‘Real-time 6DOF pose relocalization for event cameras with stacked spatial LSTM networks’. In: *Proc. IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. Workshops (CVPR)*. 2019, pp. 1638–1645.
- [17] Xavier Lagorce, Garrick Orchard, Francesco Galluppi, Bertram E. Shi and Ryad B. Benosman. ‘HOTS: A Hierarchy of Event-Based Time-Surfaces for Pattern Recognition’. In: *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)* 39.7 (2017), pp. 1346–1359.
- [18] Paul K. J. Park, Baek Hwan Cho, Jin Man Park, Kyoobin Lee, Ha Young Kim, Hyo Ah Kang, Hyun Goo Lee, Jooyeon Woo, Yohan Roh, Won Jo Lee, Chang-woo Shin, Qiang Wang and Hyunsurk Ryu. ‘Performance improvement of deep learning based gesture recognition using spatiotemporal demosaicing technique’. In: *IEEE Int. Conf. Image Process. (ICIP)*. IEEE, Sept. 2016, pp. 1624–1628.
- [19] Olaf Ronneberger, Philipp Fischer and Thomas Brox. ‘U-Net: Convolutional Networks for Biomedical Image Segmentation’. In: *Proc. IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*. May 2015.
- [20] Deqing Sun, Stefan Roth and Michael J. Black. ‘A quantitative analysis of current practices in optical flow estimation and the principles behind them’. In: *Int. J. Comput. Vis. (IJCV)* 106.2 (2014), pp. 115–137.