Multi-Event-Camera Depth Estimation and Outlier Rejection by Refocused Events Fusion

Suman Ghosh and Guillermo Gallego

Abstract—Event cameras are bio-inspired sensors that offer advantages over traditional cameras. They operate asynchronously, sampling the scene at microsecond resolution and producing a stream of brightness changes. This unconventional output has sparked novel computer vision methods to unlock the camera's potential. Here, the problem of event-based stereo 3D reconstruction for SLAM is considered. Most event-based stereo methods attempt to exploit the high temporal resolution of the camera and the simultaneity of events across cameras to establish matches and estimate depth. By contrast, this work investigates how to estimate depth without explicit data association by fusing Disparity Space Images (DSIs) originated in efficient monocular methods. Fusion theory is developed and applied to design multi-camera 3D reconstruction algorithms that produce state-of-the-art results, as confirmed by comparisons with four baseline methods and tests on a variety of available datasets.

VIDEO AND CODE

Project page: https://github.com/tub-rip/dvs_mcemvs

1 INTRODUCTION

Intelligent navigation in our complex 3D world relies on robust and efficient visual perception, which is challenging for autonomous robots. However, humans use vision very efficiently to navigate 3D environments, even in novel scenarios. Inspired by human vision, neuromorphic spikebased sensing and processing has been recently investigated for robot vision [1], [2] and retinal implants [3].

Event cameras, such as the Dynamic Vision Sensor [4]– [6] (DVS), are neuromorphic sensors that acquire visual information very differently from traditional cameras. They sample the scene asynchronously, producing a stream of spikes, called "events", that encode the time, location and sign of per-pixel brightness changes. Event cameras possess outstanding properties compared to traditional cameras: very high dynamic range (HDR), high temporal resolution ($\approx \mu$ s), temporal redundancy suppression and low power consumption. These properties offer potential to tackle challenging scenarios for standard cameras (high speed and/or HDR) [7]–[11]. However, this calls for novel methods to process the unconventional output of event cameras in order to unlock their capabilities [2].

In this work, we tackle the problem of event-based stereo 3D reconstruction for Visual Odometry (VO) and Simultaneous Localization and Mapping (SLAM). An efficient SLAM system is critical for the navigation of autonomous intelligent agents (like field robots) in challenging unstructured environments, especially in extreme ones like offshore drilling, nuclear power plants, etc. [12]. An event-based SLAM system has the potential to efficiently



Fig. 1: Semi-dense depth maps estimated by event-based monocular [10] and stereo methods. Stereo is beneficial for more accurate estimation and outlier removal compared to monocular depth estimation (e.g., traffic sign in the center). Depth is pseudo-colored, from red (close) to blue (far). Color frames are only shown for visualization. Data from [16].

overcome difficult scenarios in these applications [13]–[15]. For example the HDR advantages of event cameras translate into high-fidelity depth maps in difficult lighting conditions, as demonstrated by a broad variety of works in VO/SLAM [13, Fig. 12], [14, Fig. 8], [15, Fig. 15].

Our work is inspired by EVO [13], which is the state of the art in event-based monocular VO. The effectiveness of EVO is largely due to its mapping module, Event-based Multi-View Stereo (EMVS) [10], which enables 3D reconstruction without the need to recover image intensity, without having to explicitly solve for data association between events, and without the need of a GPU (it is fast on a standard CPU -e.g., speed of 1.20 Mev/s/core [10]). Additionally, EMVS admits an interpretation in terms of event refocusing or event alignment (contrast maximization) [17], which is the state of the art framework to tackle other vision problems [18]-[26]. Our goal is to extend EMVS to the multi-camera setting (i.e., two or more event cameras in a multi-view configuration sharing a common clock), and in particular to the stereo setting, in order to benefit from these advantages and connections (Figure 1). In the process, we revisit the event simultaneity assumption used in stereo depth estimation and develop a theory of fusion of

[•] S. Ghosh and G. Gallego are with the Technische Universität Berlin, Berlin, Germany. G. Gallego is with the Einstein Center Digital Future and the Science of Intelligence Excellence Cluster, Berlin, Germany.

Preprint of paper accepted at Advanced Intelligent Systems, 2022. doi: 10.1002/aisy.202200221.

refocused events, which could be useful in other problems, such as feature or camera tracking [27].

In summary, our contributions are:

- Simple, efficient and extensible solutions to the problem of event-based stereo 3D reconstruction for SLAM using a correspondence-free approach. We investigate early event data fusion strategies in *two orthogonal directions*: between cameras ("spatial stereo", Section 3.3) and along time ("temporal stereo", Section 3.5).
- The investigation of several functions to fuse refocused events (using e.g., Generalized means, Section 3.4) and its application to the two mentioned directions.
- A comprehensive experimental evaluation on five publicly available datasets and comparing against several baseline methods, producing state-of-the-art results (Section 4). We also show how the method can naturally handle multi-event-camera setups with *linear* complexity.

This research aims at developing robust multi-camera visual perception systems for the navigation of artificial intelligent systems in challenging environments, like stereo depth perception for SLAM and attention in robots [28], [29].

2 RELATED WORK

What: Stereo depth estimation using event cameras has been an interesting problem ever since the first event camera was invented by Mahowald and Mead in the 1990s [30]. As such, they simultaneously designed a stereo chip [31] to implement Marr and Poggio's cooperative stereo algorithm [32]. This approach has inspired a lot of literature that focuses on 3D reconstruction over short time intervals ("instantaneous stereo") [33]–[35]. These methods work well with stationary cameras in uncluttered scenes (where events are caused only by few moving objects), thus enabling 3D reconstruction of sparse, dynamic scenes. For a detailed survey on these stereo methods we refer to [36], [37]. In contrast, stereo event-based 3D reconstruction for VO/SLAM has been addressed recently [15], [38]. It assumes a static world and known camera motion (e.g., from a tracking method) to assimilate events over longer time intervals, so as to increase parallax and produce more accurate semidense depth maps. Some other works estimate depth by combining an event camera with other devices, such as light projectors [8], [39], [40] or a motorized focal lens [41], which are different from our hardware setup and application.

How: Depth estimation with stereo event cameras is predominantly based on exploiting the epipolar constraint and the assumption of *temporal coincidence* of events across retinas, namely that a moving object produces events of same timestamps on both cameras [39], [42], [43]. This aims at exploiting the high temporal resolution and redundancy suppression of event cameras to establish event matches across image planes and then triangulate. It is also known as *event simultaneity* or *temporal consistency* [38], and it is analogous to photometric consistency in traditional cameras. This assumption does not strictly hold [44], [45], and so it is relaxed to account for temporal noise (jitter and delay). Essentially event simultaneity is exploited to solve the data association problem (establishing event matches), which is



Fig. 2: *DSI projections* of a non-planar scene (*rpg_monitor* data from [38], also in Figure 5). Disparity Space Image (DSI) values (i.e., ray counts) are pseudo-colored, from blue (low) to yellow (high). The DSI has dimensions $w \times h \times N_Z$, according to the resolution of the event camera DAVIS240 (w = 240, h = 180 pix) and the number of *inverse-depth* planes used, $N_Z = 100$. The three DSI max-projections are: front view (top left, 240×180 pix), top view (bottom left, 240×100 pix) and side view (top right, 100×180 pix).

a well-known difficult problem due to the little information carried by each event and their dependency with motion direction (changing "appearance" of events [2], [46]).

The above ideas are used in the mapping module of [15], the state-of-the-art stereo 3D reconstruction method for VO/SLAM. In this method temporal consistency is measured across space-time neighborhoods of events by first converting the events into time surfaces (TSs) [47] and then comparing their spatial neighborhoods. Stereo point matches are established and provide depth estimates which are fused in a probabilistic way using multiple TSs to produce a more accurate semi-dense inverse depth map.

In contrast, we investigate a new way of doing stereo, without explicitly using event simultaneity and hence without establishing event matches. Therefore, to the best of our knowledge, we completely depart from previous eventbased stereo methods. Correspondence-free approaches for depth estimation have been proposed for frame-based cameras [48]-[50]. For example, [50] proposes a method to estimate affine fundamental matrices without explicit correspondences, but it does not solve the full problem of depth estimation. On the other hand, [49] aims to solve the problem of stereo depth estimation, but for the special case of planar or quasi-planar scenes. It is an extension of [51] for arbitrary camera positions. In contrast, our method solves the full depth estimation problem and is not limited to the planar case, following the seminal idea of [48] to sweep space and exploit the sparsity of scene edges. Inspired by [10], [48], we circumvent the data association task by leveraging the sparsity of events (event cameras naturally highlight edges, which are sparse, in hardware) and by exploiting the continuous set of camera viewpoints at which events are available. This provides a rich collection of back-projected rays through the events to estimate scene structure. Our contributions pertain to the processing (e.g., fusion) of such back-projected rays or "refocused events", which has not been considered before (since [10] and newer approaches [52] do not consider data fusion, e.g., across cameras).



Fig. 3: Our method takes as input the events from two or more synchronized, rigidly attached event cameras and their poses, and estimates the scene depth. Using Space Sweeping, it builds ray density Disparity Space Images (DSIs) from each camera data and fuses them into one DSI (Section 3.4), from which the 3D structure of the scene is extracted in the form of a semi-dense depth map, which may be cast into a point could. Fusion across cameras (Section 3.3) is represented with yellow lines, and temporal fusion (Section 3.5) with red lines. The optional shuffling block ("S") is presented in Section 3.6.

3 EVENT-BASED STEREO DEPTH ESTIMATION

This section reviews how an event camera works (Section 3.1) and the monocular method EMVS (Section 3.2) before presenting our stereo depth estimation approach. Two main event fusion directions are presented: fusion of camera views (Section 3.3) using one of several functions (Section 3.4), and fusion of multiple time intervals (Section 3.5). Then, we revisit the event simultaneity assumption (Section 3.6) and analyze the computational complexity of the approach (Section 3.7).

3.1 How an Event Camera Works

Event cameras, such as the Dynamic Vision Sensor (DVS) [4], are bio-inspired sensors that capture pixel-wise *brightness changes*, called events, instead of brightness images. An event $e_k \doteq (\mathbf{x}_k, t_k, p_k)$ is triggered when the logarithmic brightness *L* at a pixel exceeds a contrast sensitivity $\theta > 0$,

$$L(\mathbf{x}_k, t_k) - L(\mathbf{x}_k, t_k - \Delta t_k) = p_k \,\theta,\tag{1}$$

where $\mathbf{x}_k \doteq (x_k, y_k)^{\top}$, t_k (in µs) and $p_k \in \{+1, -1\}$ are the spatio-temporal coordinates and polarity of the brightness change, respectively, and $t_k - \Delta t_k$ is the time of the previous event at the same pixel \mathbf{x}_k . Hence, each pixel has its own sampling rate, which depends on the visual input. Assuming constant illumination, pixels produce events proportionally to the amount of scene motion and texture.

3.2 EMVS: Monocular 3D Reconstruction

The problem of *monocular* depth estimation with an event camera consists of estimating the 3D structure of the scene given the events and the camera poses (i.e., position and orientation) as the sensor moves through the scene. The method in [10] solves this problem, called EMVS, in two main steps: it builds a Disparity Space Image (DSI) using a space sweeping approach [48] and then detects local maxima of the DSI. The key idea is that, as the camera moves, events are triggered at an almost continuous set of viewpoints, which are used to back-project events into space in the form of rays (a DSI). The local maxima of the ray

Algorithm 1 Stereo event fusion across cameras

- 1: *Input*: stereo events in interval [0, *T*], camera trajectory, camera calibration (intrinsic and extrinsic).
- 2: Define a single reference view (RV) for both DSIs, coinciding with the left camera pose at say t = T/2.
- 3: Create 2 DSIs by back-projecting events from each camera.
- 4: Fusion: compute the pointwise harmonic mean of the DSIs.
- 5: Extract depth and confidence maps from the fused DSI f: $Z^*(x, y) \doteq \arg \max f(\mathbf{X}(x, y)), c^*(x, y) \doteq \max f(\mathbf{X}(x, y)).$

density (where many rays intersect, as shown in Figure 2) are candidate locations for the 3D edges that produce the events. Specifically, the DSI is discretized on a projective voxel grid defined at a reference view, and local maxima are detected along viewing rays, thus producing a semi-dense depth map. Events are processed in packets of about 0.20–1.00 M events. The key benefits of EMVS are its simplicity, accuracy, efficiency (real-time, with ~1.2Mev/s throughput per CPU core [10]) and that it estimates depth without explicit data association.

3.3 Fusion Across Cameras

We consider the problem of depth estimation from two synchronized and calibrated event cameras rigidly attached. Hence, the input consists of a stereo event stream and poses, and the desired output is a (semi-dense) depth map or, equivalently, a 3D point cloud with the scene structure (Figure 3).

3.3.1 Challenges and Proposed Architecture

A naive solution to the problem consists of running two instances of EMVS, one per camera, and fusing the resulting point clouds into a single one, including post-processing to mitigate redundant 3D points. This is a late-fusion approach, which greatly ignores the benefits that arise from having two cameras observing the same scene. By contrast, we seek to perform fusion earlier in the processing pipeline: at the DSI stage. Hence the first technical challenge is to define the DSI. EMVS defines a DSI per camera, located at a



(a) Event cameras and scene. 0.05 s 0.20 s 1.00 s 0.00 s 0.

(b) Evolution of DSI projections.

Fig. 4: *Intuitive Example: Monocular vs. Stereo* method on planar rock scene and 1D motion along the camera's *X* axis. Plots of the evolution of the DSI projections (see text) for different methods (rows) as time increases (columns). The 3D edge patterns in the DSI (in yellow) are less localized in EMVS (top row) than in stereo Alg. 1 (bottom).

reference view (RV) along the camera's trajectory. However, the fusion of DSIs at two different RVs is prone to resampling errors. Thus it is key to define a common DSI location for both cameras.

The second challenge is to investigate sensible fusion strategies. Our approach includes, as a particular case, that of back-projecting the events from both cameras into a single DSI and simply counting rays. This is equivalent to the scenario of a single event camera that moves twice through the scene, with different motions, but uses the same DSI to aggregate rays. It doubles the ray count in the DSI, but summation discards valuable information for fusion, such as how many rays originate in each camera: given 8 rays at a point, it is preferable to have 4 rays from each camera than an unbalanced situation (a 3D edge seen only by one camera).

To deal with the above challenges we define two DSIs at a common reference view: having one DSI per camera allows us to preserve the origin of the event data, and having *geometrically aligned* DSIs avoids resampling errors during fusion. Without loss of generality, let the RV be a point along the trajectory of the left camera. We investigate how to compare and fuse the ray densities from each event camera. Figure 3 shows the block diagram of our stereo approach. For now, assume there is $N_s = 1$ DSI per camera $(N_s > 1)$ is presented in Section 3.5). First, the aligned DSIs are populated with back-projected events from each camera, then they are *fused (combined)* into a single one (e.g., using a voxel-wise harmonic mean or other similarity



Fig. 5: *Output of Stereo Alg.* **1** *on two scenes*. Top: our method produces a semi-dense depth map of the scene (color coded from red (close) to blue (far), overlaid on a grayscale frame of the DAVIS [53]), and a confidence map (Bottom) with the maximum DSI value along each reference view pixel, in negated scale (bright = small; dark = large).

score (Section 3.4)), and finally local maxima are extracted to produce a semi-dense depth map. The steps are specified in Alg. 1.

3.3.2 Intuitive Example

To illustrate key differences between EMVS (monocular) and the stereo Alg. 1 we use a sequence acquired with two event cameras [53] performing a 1D motion (translation along the X axis, using a linear slider). As input to EMVS we use the data from the left camera. Figure 4 shows the evolution of the DSIs as time progresses, i.e., as the camera rig moves and more events are acquired and backprojected onto the DSIs (and fused in the stereo case). We plot projections of the DSI along its three coordinate axes (like in Figure 2) at the reference viewpoint RV.

As Figure 4 shows, the stereo DSI (bottom row) converges faster to the 3D structure of the scene than the monocular DSI (top row). This is specially noticeable in the top views: only one set of nearly parallel rays, poor for triangulation, is visible in the monocular case. By contrast, the top views of the stereo DSIs show two sets of rays, one from each camera: the rays from the left camera are nearly straight, whereas the rays from the right camera are curved due to inverse depth parametrization of the DSI grid and the fact that the DSI is projective. In both scenarios, the rays intersect at multiple voxels and as time progresses the true intersection locations dominate over others, i.e., the 3D structure emerges by event refocusing [10], [18]. Additionally, in the stereo case refocusing is combined with the proposed fusion functions (Section 3.4) to speed-up the emergence and better highlight 3D structure. In Alg. 1 this is achieved by the harmonic mean, which deemphasizes the non-intersecting parts of the rays.

3.3.3 Output of the Stereo Method

Figure 5 shows the output of Alg. 1 on two scenes. After DSI fusion, Alg. 1 extracts a depth map by locating the DSI max-

ima along each viewing ray (through RV pixel $\mathbf{x} = (x, y)^{\top}$). Letting $f : \mathbb{R}^3 \to \mathbb{R}_{\geq 0}$ be the fused DSI, its maxima provide the confidence or "contrast" map $c(\mathbf{x}) = f(\mathbf{x}, Z^*(\mathbf{x}))$ and the depth map $Z^*(\mathbf{x})$. Adaptive Gaussian Thresholding (AGT) selects the pixels with highest local value, thus making the depth maps semi-dense. A median filter is applied to remove isolated points. The front-view projection of the DSI in Figure 2 corresponds to the confidence map, which is called this way because it is used in AGT to "select the most confident pixels in the depth map" [10] (since voxels with many ray intersections are more likely to capture true 3D points than voxels with few ray intersections).

3.4 DSI Fusion Functions

DSI fusion is the central part of our method (Figure 3). It takes two ray density DSIs on the same region of space as input (one per camera) and produces a merged DSI, which is then used to extract depth information (candidate locations of 3D edges). So, *what are sensible ways to compare two DSIs?* FMRay density DSIs have very different statistics from natural images, hence standard similarity metrics for image patches may not be the most appropriate ones [54].

Formally, let $\mathcal{E}_l = \{e_k^l\}_{k=1}^{N_e^l}$ and $\mathcal{E}_r = \{e_k^r\}_{k=1}^{N_e^l}$ be stereo events over some time interval [0, T], and $f_l, f_r : V \subset \mathbb{R}^3 \to \mathbb{R}_{\geq 0}$ be the ray densities (DSIs) defined over a volume V:

$$f_l(\mathbf{X}) = \sum_{k=1}^{N_e^l} \delta\big(\mathbf{X} - \mathbf{X}'_k(e_k^l)\big),\tag{2}$$

where $\mathbf{X}'_k(e_k^l) = (\mathbf{x}_k^{l/\top}, Z)^\top$ is a 3D point on the backprojected ray through event e_k^l , at depth Z with respect to the reference view (RV). Events are transferred to RV using the continuous motion of the cameras and candidate depth values $Z \in [Z_{\min}, Z_{\max}]$:

$$\mathbf{x}_{k}^{l\prime} = \mathbf{W}\big(e_{k}^{l}, \mathsf{P}^{l}(t_{k}), \mathsf{P}_{v}, Z\big),\tag{3}$$

where $P^{l}(t)$ is the pose of the left event camera at time t and P_{v} is the pose of RV. The warp **W** corresponds to the planar homography induced by a plane parallel to the image plane of RV and at the given depth Z. In a coordinate system adapted to RV (i.e., $P_{v} = (I, 0)$), the planar homography is given by the 3×3 homogeneous matrix

$$\mathbf{H}_{\mathbf{W}} \sim \left(\mathbf{R} + \frac{1}{Z} \mathbf{t} \mathbf{e}_{3}^{\top}\right)^{-1},\tag{4}$$

where $P^l(t_k) = (\mathbf{R}, \mathbf{t})$ and $\mathbf{e}_3 = (0, 0, 1)^{\top}$. A similar formula applies to compute f_r from \mathcal{E}_r and the corresponding camera poses. In practice, DSIs are discretized over a projective voxel grid with N_Z depth planes in $[Z_{\min}, Z_{\max}]$, and the Delta δ in (2) is approximated by bilinear voting [10], [55]. Hence, each voxel counts the number of event rays that pass through it.

Next, letting $u = f_l(\mathbf{X})$, $v = f_r(\mathbf{X})$ be the values of the DSIs at a 3D point \mathbf{X} , we seek to define a fused value $g(\mathbf{X})$. For simplicity, we consider metrics operating in a point-wise (i.e., voxel-wise) manner, i.e., with a slight abuse of notation:

$$g(\mathbf{X}) \equiv g(f_l(\mathbf{X}), f_r(\mathbf{X})) = g(u, v).$$
(5)



Fig. 6: Fusion functions considered, for u = [0, 4], v = 1.

The metrics considered are the following:

$A(u,v) \doteq (u+v)/2$	Arithmetic mean	(6)
$G(u,v) \doteq \sqrt{uv}$	Geometric mean	(7)
$H(u,v) \doteq 2/(u^{-1} + v^{-1})$	Harmonic mean	(8)
$\operatorname{RMS}(u,v) \doteq \sqrt{\frac{1}{2}(u^2 + v^2)}$	Quadratic mean	(9)
$\min(u, v)$	Minimum	(10)
$\max(u, v)$	Maximum	(11)

They are special cases of the Generalized mean (power mean or Hölder mean) and satisfy an order (see Figure 6):

$$\min \le H \le G \le A \le \text{RMS} \le \max,\tag{12}$$

where the equal sign holds if and only if u = v. Eq. (12) also establishes a qualitative order of the depth maps obtained after DSI fusion with the corresponding function. The arithmetic mean A (i.e., averaging ray densities) corresponds to the above-mentioned particular case of counting the backprojected stereo events on a single DSI. Hence, functions performing worse than this case are not pursued.

3.4.1 What makes a good fusion function?

Intuitively, given two ray densities defined on the same volume, a fusion function should emphasize the regions of high ray density on both DSIs and deemphasize the rest. It is not sufficient for one of the two densities to be large at a point **X** to signal the presence of a 3D edge; *both* densities have to be similar and large at **X**. The arithmetic mean *A* and its dominant functions (e.g., quadratic mean and max in (12)) do not satisfy this "AND" logic, whereas the geometric mean, harmonic mean and min functions do satisfy it (e.g., a large value G(u, v) can only be achieved if both u and v are large).

Mathematically, this requirement is well described by the *concavity* properties of the function (plot in Figure 6). The arithmetic mean A and functions below it are concave (assuming non-negative inputs). Further, G, H and min are strictly concave. As the experiments will show, fusion using G still lets notorious outliers pass. Functions H and min deemphasize considerably more than G. The min function saturates strictly, treating values u > v as u = v, hence clipping and discarding potentially beneficial information about ray density values. The harmonic mean H shows strong concavity and varies smoothly with both input arguments, without discarding information. H is dominated by the minimum of its arguments,

$$\min(u, v) \le H(u, v) \le 2\min(u, v) \tag{13}$$

(in terms of the plot in Figure 6 (v = 1), the green curve is bounded: $H(u, 1) \leq 2$). The goal of the present work is to introduce and study fusion functions rather than to select a single "best one". To narrow the discussion we often use a subset of the fusion functions.

3.4.2 Interpretation in terms of Contrast Maximization

The proposed stereo fusion method is related to contrast/focus maximization [17], [18]. The depth slices of the DSIs count refocused events (warped by back-projection), i.e., they constitute so-called images of warped events (IWEs) [17]. The fused DSI can be interpreted as a similarity score between refocused events. Since fusion functions such as H try to emphasize DSI regions with large and similar values, stereo Alg. 1 tries to maximize the similarity score between refocused events (in-focus effect) on both cameras, jointly. The confidence map registers the maximum focus similarity score at each viewing ray of the fused DSI.

3.4.3 More fusion functions

Additional means exist beyond those in Figure 6, such as the contraharmonic mean, logarithmic mean, quasi-arithmetic mean, arithmetic-geometric mean, Heronian mean and weighted generalized means. However they are not covered for the sake of brevity and to avoid clutter. In some cases, the order relation (12) can be extended to justify their limited practical interest. Seeking more fusion functions, one could combine functions in Figure 6 with non-linear transformations of the input DSIs, in a homomorphic filtering fashion. For example, the A-mean of the log-DSIs is related to the *G*-mean of the DSIs, which has a stronger concavity than the A-mean of the DSIs. The same idea can be applied to other functions to increase concavity: the G-mean of the log-DSIs, $G(\log(1+u), \log(1+v)))$, has stronger concavity than the G-mean of the original DSIs. The logarithm plays down large DSI values, thus deemphasizing differences between corresponding DSIs before fusion. For simplicity, we restrict the study to the functions in Figure 6.

3.4.4 Loose connection with prior fusion work

A method for fusing 3D representations called "temporally synchronized event disparity volumes" was proposed in [56], where two binary data volumes I_L , I_R were fused using an intersection-over-union (IoU) cost. It resembles the H mean:

$$IoU = \frac{\sum_{\mathbf{x}\in W} I_L(\mathbf{x}, d) \cap I_R(\mathbf{x}, d)}{\sum_{\mathbf{x}\in W} I_L(\mathbf{x}, d) \cup I_R(\mathbf{x}, d)} \text{ vs. } H = 2\frac{uv}{u+v}, \quad (14)$$

where the product in the numerator ("intersection") acts as an "AND" condition and the sum in the denominator ("union") acts as a normalization factor. However, note that the IoU in [56] is computed by aggregating binary data I_L , I_R over spatial windows W (of 32×32 pixels), whereas H (Figure 6) is computed voxel-wise, without spatial aggregation (i.e., it has higher spatial resolution), on continuous ray densities.

Algorithm 2 Stereo event fusion across cameras and time

- 1: *Input*: stereo events in interval [0, *T*], camera trajectory, camera calibration (intrinsic and extrinsic).
- 2: Define a single reference view (RV) for all DSIs.
- 3: Divide the interval [0, T] into N_s sub-intervals (of equal size or equal number of events). Create $2N_s$ DSIs by backprojecting events from each subinterval and camera.
- 4: $S_2 \circ S_1$: Two fusion axes (cameras and time). If $S_2 \equiv A_t$ and $S_1 \equiv H_c$, compute first the S_1 fusion (*H*-mean of two corresponding DSIs, on the same sub-interval); then compute the S_2 fusion (*A*-mean of all sub-interval DSIs).
- 5: Extract depth and confidence maps from the fused DSI.

3.5 Temporal Fusion

The functions presented in Section 3.4 can be used to fuse any pair of aligned DSIs. Moreover, the functions can be extended to handle more than two inputs: they allow us to *fuse an arbitrary number of registered DSIs*, with a complexity that is *linear* in the number of DSIs. The DSIs may be populated by events from different cameras or, as we also investigate, from different time intervals. The main idea is to split an interval into multiple sub-intervals, build the DSI for each of them and fuse all DSIs into a single one (Figure 3). This strategy can be applied regardless of the number of cameras in the system, hence it represents an independent axis of variation. Moreover, *the same technique enables cameraand time- fusion*, which we collectively call Alg. 2. The key lines of Alg. 2 that change with respect to Alg. 1 are lines 3 and 4.

Given *N* fusion functions (N = 6 in Figure 6) there are $2N^2$ possible fusion schemes considering the choice of temporal fusion function, across-camera fusion function and the order of application (Alg. 2). For brevity we reduce the analysis to the comparison of N = 2 fusion functions: *A* and *H*, which yield 8 possible fusion schemes.

Let A_t denote the fusion operation along the time (t) axis using the arithmetic mean (A). Likewise, H_c is the fusion operation along the camera (c) axis using the harmonic mean (H). Then, $A_t \circ H_c$ first applies H_c (producing as many DSIs as sub-intervals) and then A_t . Out the of 8 possibilities, there are only 6 distinct ones due to commutativity:

$$A_c \circ A_t = A_t \circ A_c$$
 and $H_c \circ H_t = H_t \circ H_c$. (15)

The four remaining fusion combinations are:

$$A_t \circ H_c, A_c \circ H_t, H_t \circ A_c \text{ and } H_c \circ A_t.$$
 (16)

Clearly, $A_c \circ A_t$ is equivalent to the approach of summing all stereo events into a single DSI, and $H_t \circ H_c$ is very restrictive because only edges seen by all cameras in all subintervals will survive. Alg. 1 is also a particular case of Alg. 2 ($H_c \circ A_t$).

3.6 Is Event Simultaneity Needed in Stereo?

Data association is a fundamental problem in event-based vision [2]. In stereo, event simultaneity is a cornerstone assumption to resolve data association (i.e., find corresponding points) and subsequently infer depth. A thought-provoking discovery made while developing our method is that across-camera fusion does not need to be done on corresponding intervals (step 4 in Alg. 2). We tested our method with the shuffling block in Figure 3 enabled and it

still produced good results (see Section 4.4.3). Hence, the proposed stereo method foregoes the event simultaneity assumption. The explanation is that given the camera poses, events are transformed into a representation (i.e., the DSI) where event simultaneity is not as critical as in the instantaneous stereo problem (Section 2). The camera poses serve as a proxy allowing us to reproject stereo events to a common DSI and fuse them, even if the DSIs are well separated in time. The DSI representation is sufficient to produce 3D reconstructions. Stereo is not solved by matching events, but by comparing possibly non-simultaneous DSIs (each DSI spans several thousands of events).

3.7 Complexity Analysis

Let us analyze the complexity of the proposed stereo methods in comparison with the monocular case. The main steps of the methods are: DSI creation (event back-projection), DSI fusion, maxima detection along viewing rays of the DSI, and thresholding (AGT). If N_e is the number of events, N_p is the number of pixels in the reference view, N_Z is the number of depth planes in the DSI, and N_k is the number of pixels in the AGT kernel (e.g., 5×5), then the complexity of [10] is

$$O(\underbrace{N_e N_Z}_{\text{DSI creation}} + \underbrace{N_Z N_p}_{\text{arg max}} + \underbrace{N_p N_k}_{\text{AGT}}).$$
(17)

In the case of Alg. 1 with N_c cameras, assuming that each camera produces N_e events, there are N_c DSIs to build and fuse. Hence, the complexity is:

$$O(\underbrace{N_c N_e N_Z}_{\text{DSI creation}} + \underbrace{N_c N_Z N_p}_{\text{DSI fusion}} + \underbrace{N_Z N_p}_{\text{arg max}} + \underbrace{N_p N_k}_{\text{AGT}}).$$
(18)

In the case of Alg. 2 with N_s subintervals, there are N_s DSIs per camera, but each one has N_e/N_s events, and so the complexity of DSI creation does not change. Only the fusion step becomes more expensive:

$$O(\underbrace{N_c N_e N_Z}_{\text{DSI creation}} + \underbrace{N_s N_c N_Z N_p}_{\text{DSI fusion}} + \underbrace{N_Z N_p}_{\text{arg max}} + \underbrace{N_p N_k}_{\text{AGT}}).$$
(19)

4 EXPERIMENTS

To assess the performance of our method we test on a wide variety of real-world and synthetic sequences, which are introduced in Section 4.1. Section 4.2 compares functions for fusion across cameras. Section 4.3 compares our method with three state-of-the-art methods on MVSEC and UZH data. Section 4.4 evaluates temporal fusion and sub-interval shuffling. Then, we evaluate on higher resolution data: driving dataset DSEC (Section 4.5), 1Mpixel VIO dataset TUM-VIE (Section 4.6), and analyze the sensitivity with respect to the camera's spatial resolution (Section 4.7). We also present trinocular examples (Section 4.8), analyze runtime (Section 4.9) and sensitivity with respect to the camera's contrast threshold (Section 4.10). Finally, Section 4.11 summarizes the findings and Section 4.12 discusses limitations of the method.

TABLE 1: Parameters of stereo or trinocular event-camera rigs used in the experiments.

Dataset	Cameras	Resolution [pix]	Baseline [cm]	FOV [°]
ECCV18 [38]	DAVIS240C	240×180	14.7	62.9
MVSEC [57]	DAVIS346	346×260	10.0	74.8
EVIMO2 [58]	Samsung Gen3	640×480	trinocular	75
	2× Prophesee Gen3	640×480	trinocular	70
DSEC [16]	Prophesee Gen3	640×480	60	60.1
TUM-VIE [59]	Prophesee Gen4	1280×720	11.84	90
ESIM [60]	Simulator	up to 1280×960	20	77.3

4.1 Datasets and Evaluation Metrics

4.1.1 Datasets

We evaluate our stereo methods on sequences from five publicly available datasets [16], [38], [57]–[59] and a simulator. Sequences from [38], [58] were acquired with a hand-held stereo or trinocular event camera in indoor environments. Sequences in the MVSEC dataset [57] were acquired with a stereo event camera mounted on a drone while flying indoors. The sequences in the DSEC dataset [16] were recorded with event cameras on a car that drove through Zurich's surroundings. The TUM-VIE dataset was recorded with the sensor rig mounted on a helmet, and its sequences contain indoor and outdoor scenes. The simulator [60], [61] provides synthetic sequences using an ideal event camera model and scenes built using CAD models.

Ground Truth. Some datasets contain ground truth *poses* from a motion-capture system, which we use as input to all tested methods. If camera poses are not available (e.g., TUM-VIE), we compute them using data from the sensor rig (e.g., a visual-inertial odometry algorithm). Some datasets, such as MVSEC and DSEC, contain ground truth *depth* for quantitative assessment of the 3D reconstruction methods. Depth is given by a LiDAR operating at 10.00–20.00 Hz. The event camera pixels corresponding to points outside the LiDAR's field of view (FOV) or points close to the sensor rig may not have a LiDAR depth value.

Rigs and Calibration. The main geometric parameters of the event cameras used in the above datasets are summarized in Table 1. The stereo rigs in [38], [57] consist of two Dynamic and Active Pixel Vision Sensors (DAVIS) [53]. The DAVIS comprises a frame-based and an event-based sensor on the same pixel array, thus calibration (intrinsic and extrinsic) is achieved using the intensity frames, and then it is applied to the events. The datasets whose cameras output only events (EVIMO2, DSEC and TUM-VIE), are calibrated by converting events to frames and calibrating the latter (e.g., using [62]). All methods work on undistorted coordinates.

4.1.2 Metrics

The performance of the proposed method is quantitatively characterized using several standard metrics on the datasets with ground truth depth (i.e., MVSEC and DSEC). We provide mean and median errors between the estimated depth and the ground truth one (median errors are more robust to outliers than mean errors). We also report the number of reconstructed points, the number of outliers (bad-pix [63]), the scale invariant depth error (SILog Err), the sum of absolute value of relative differences in depth (AErrR), and



Fig. 7: *Fusion Functions*. Semi-dense depth maps (top rows) and confidence maps (bottom row) produced by Alg. 1 using the fusion functions in Figure 6 on data from [38]. The columns (fusion functions) follow the order in (12). The differences are subtle: the depth maps on the left columns have fewer outliers than those on the right columns (zoomed-in insets). The differences are more noticeable in the confidence maps, whose sharpness increase from right to left. Depth is color coded, from red (close) to blue (far), in the range 0.55–6.25 m. Confidence maps are colored as in Figure 5.

TABLE 2: Depth errors for stereo DSI fusion across cameras (Section 4.2). Experiments on 200 s (110 million events) of the three indoor flying sequences from MVSEC [57]. The maximum scene depth is 8.40 m.

	Mean A	Mean Abs Error [cm] \downarrow			Median Abs Error [cm] \downarrow			bad-pix [%] \downarrow			#Points [million] ↑		
Sequence [57]	flying1	flying2	flying3	flying1	flying2	flying3	flying1	flying2	flying3	flying1	flying2	flying3	
$\min_c \circ A_t$	58.13	68.79	49.96	24.07	39.20	20.35	17.40	38.82	12.40	6.13	12.72	5.83	
Alg. 1 ($H_c \circ A_t$)	60.16	68.83	51.94	25.45	39.63	21.15	18.58	38.91	13.70	6.61	13.54	6.27	
$G_c \circ A_t$	63.57	70.48	55.22	28.17	41.51	22.97	20.30	39.68	15.30	6.90	14.18	6.35	
$A_c \circ A_t$	78.78	79.95	60.35	38.15	47.63	25.03	27.20	44.64	17.56	5.97	14.53	4.50	
$RMS_c \circ A_t$	99.15	88.46	88.21	61.73	55.71	47.45	36.84	48.89	29.86	7.65	17.48	5.41	
$\max_c \circ A_t$	109.30	93.81	104.52	75.44	61.71	66.56	41.31	51.64	36.27	9.25	19.99	6.92	

 δ -accuracy values on the percentage of points whose depth ratio with respect to ground truth is within some threshold (see [64]). We also provide precision, recall and F1-score curves [65]. *Precision* is the percentage of estimations that are within a certain error from the ground truth. *Recall* (e.g., completeness or reconstruction density) is the percentage of ground truth points that are within a certain error from the estimations. The *F1 score* is the harmonic mean of precision and recall, which is dominated by the smallest of them. Since the depth maps obtained are semi-dense (while the ground truth is often more dense), recall often dominates.

The method in Section 3 is presented for the events in a time window. To apply the method to a whole sequence, we split the latter into non-overlapping time windows and apply the method to each of them. When thresholding to obtain semi-dense depth maps (AGT step), we normalize by a robust maximum DSI value obtained over the sequence, which makes the comparisons more stable.

4.2 Comparison of Across-Camera Fusion Functions

We first evaluate Alg. 1 using the fusion functions in Figure 6. Figure 7 shows qualitatively the corresponding depth- and confidence maps for a sample sequence. The columns follow the order in (12). The arithmetic mean (A) corresponds to counting the back-projected event rays from

both cameras (left/right) on a single DSI. It has not been proposed before in the literature and is a particular case of our fusion methods. It provides moderate results by conveying that a 3D point is detected if enough rays are counted at a voxel, regardless of which camera the event ray originated from. However, it does not filter out spurious ray intersections (which do not correspond to actual 3D points) until there is sufficient evidence. Spurious ray intersections are more common in the stereo case than in the monocular one because rays are originated from two moving sources in stereo instead of just one. The columns in Figure 7 to the right of A (i.e., quadratic mean and max) produce worse results than A. They convey that 3D points are detected if enough rays from at least one camera are counted at a voxel. This strategy might be good to mitigate occlusions (edges seen in only one camera), but it does not produce optimal results if the edge is visible from both event cameras. Finally, the columns to the left of A (i.e., G, H, \min) produce better results than A. This is due to the fact that they implement a more conservative strategy: 3D points are detected only if enough rays from both cameras are counted at a voxel. The results are more noticeable in the confidence maps; these are sharper in the first columns than in the last ones.

Table 2 quantitatively compares the six fusion functions on the three indoor flying sequences from [57]. The experi-

9



Fig. 8: *Event-based Stereo 3D Reconstruction*. Comparison of depth estimation results on several sequences using various stereo methods. For visualization purposes, the first column depicts intensity frames from the DAVIS camera (not used by any method). Columns 2 to 7 show semi-dense inverse depth maps produced by GTS [43], SGM [66], ESVO [15], our Alg. 1 and Alg. 2 ($A_t \circ H_c$ and $A_c \circ A_t$, with $N_s = 2$), respectively. Depth maps are pseudo-colored, from red (close) to blue (far), in the range 0.55–6.25 m for the *rpg* sequences [38] and in the range 1.00–6.50 m for the *upenn* MVSEC sequences [57].

ment consists of running stereo Alg. 1 on 200 s of data (110 million events), at 20 Hz. The estimated depth produces \approx 28.4 million points on \approx 4000 ground truth snapshots for each fusion function. The differences between fusion functions are most noticeable when less data is available, and so we use events packets of 0.1s for this experiment. Table 2 reports mean and median depth errors, bad-pixel percentage and number of reconstructed points. Errors and bad-pix follow a clear trend, decreasing towards the top rows. The number of reconstructed points also decreases, but non-monotonically in sequences 1 and 3. Median errors are considerably smaller than mean errors, signaling the presence of outliers (points with large depth errors). We observe an accuracy-completion trade-off: the top rows are more accurate than the bottom rows, but the bottom rows provide more reconstructed points. Comparing the top two rows (highest accuracy), the differences in accuracy are small (mean: 2.41%, median: 3.43%) while the difference in number of points is larger: 6.83%. Hence these results indicate, together with theoretical aspects (Section 3.4), that the *H*-mean is advantageous to fuse across cameras.

4.3 Comparison with Stereo State of the Art

We assess the performance of our methods in comparison to several event-based stereo methods, in Figure 8 and Table 3.

4.3.1 Baseline Methods

The Generalized Time-Based Stereovision method (GTS) [43] follows a classical two-step approach: stereo

matching plus triangulation. Matching is based on a per-event time-based consistency score. The Semi-Global Matching (SGM) method [66] is adapted to event data by feeding time images [47] and masking the produced depth map at recent event locations [38]. We also compare against the mapping module of Event-based Stereo Visual Odometry (ESVO) [15], which fuses multiple depth estimates using Student-t filters, with each estimate and its uncertainty obtained by maximizing spatio-temporal consistency between patches of stereo time images. GTS and SGM are also endowed with depth propagation-and-fusion filters, as implemented in [15]. Finally, we also compare stereo against the monocular method EMVS [10], which has not been carried out before. All baseline methods produce depth maps at the LiDAR rate (20 Hz) and use ground truth poses to propagate depth estimates in time, if needed. ESVO on MVSEC data works by fusing 20 depth maps generated at 20 Hz, i.e., 1 s of data. EMVS works on event packets of 1 s, shifted by 50 ms (20 Hz). For a sensible comparison with EMVS, baseline stereo methods are also run on event packets of 1 s, shifted by 50 ms; this is highlighted as "indep 1s" in Table 3.

4.3.2 Results

Figure 8 compares qualitatively the inverse depth maps produced by the above stereo methods. To illustrate the appearance of the scenes, the first column shows grayscale frames from the DAVIS [67]. The remaining columns show the output of GTS, SGM, ESVO and our methods. Because

	Algorithm	Mean Err [cm]↓	Median Err [cm]↓	bad-pix [%]↓	SILog Err ×100↓	AErrR [%]↓	$\begin{array}{c} \log \mathrm{RMSE} \\ \times 100 \downarrow \end{array}$	$\begin{array}{c} \delta < 1.25 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^2 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^3 \\ [\%] \uparrow \end{array}$	#Points [million]↑
SOTA	EMVS [10] (monocular) ESVO [15] ESVO indep. 1s SGM indep. 1s GTS indep. 1s	33.78 25.00 22.70 35.42 389.00	14.35 10.59 9.83 12.35 45.43	3.84 3.35 2.83 6.39 38.45	4.20 3.48 3.03 8.45 74.47	12.74 10.19 9.59 16.17 102.92	20.72 18.83 17.53 29.49 89.08	84.75 90.44 91.82 85.34 49.56	94.87 95.76 96.50 93.05 62.19	97.99 97.98 98.38 96.03 69.36	1.27 2.04 1.56 14.46 0.06
Ours	$ \begin{array}{l} H_c \circ A_t \ (\mathrm{Alg} \ 1) \\ H_c \circ A_t \ (\mathrm{Alg} \ 1) + \mathrm{MF} \\ H_c \circ H_t \\ H_t \circ A_c \\ A_c \circ H_t \\ A_c \circ A_t \\ A_t \circ H_c \\ A_t \circ H_c \\ A_t \circ H_c + \mathrm{shuffling} \end{array} $	20.07 20.64 23.45 22.61 23.35 <u>20.38</u> 20.92 22.60	9.53 9.72 10.98 10.68 10.94 <u>9.60</u> 9.76 10.71	$ \begin{array}{r} 1.35 \\ \underline{1.43} \\ 1.89 \\ 1.67 \\ 1.84 \\ 1.51 \\ 1.66 \\ 1.68 \\ \end{array} $	1.72 1.80 2.18 2.03 2.15 1.80 1.86 2.11	7.80 7.94 8.86 8.61 8.83 <u>7.93</u> 8.05 8.58	13.24 <u>13.54</u> 14.93 14.44 14.85 13.55 13.81 14.28	95.04 <u>94.74</u> 93.00 93.49 93.07 94.67 94.39 93.49	98.08 97.95 97.49 97.76 97.53 <u>98.01</u> 97.80 97.76	99.21 99.17 99.03 99.12 99.04 <u>99.20</u> 99.14 99.13	$\begin{array}{c} 0.81\\ \underline{3.00}\\ 1.47\\ 1.25\\ 1.42\\ 0.99\\ 1.14\\ 1.24 \end{array}$

TABLE 3: Quantitative evaluation and comparison of our proposed method with the state of the art. All metrics are averaged over the three indoor MVSEC sequences (flying 1, 2, 3), where the maximum ground truth depth is 8.40 m. The methods are evaluated on 200 s of data (110 million events and 4000 ground truth depth maps). Each estimated depth map is computed using 1 s of event data (≈ 0.55 million events). MF: morphological filter (see text). Per-sequence results are in Tables 7 to 9.

event cameras naturally respond to the apparent motion of edges, which occupy only a small portion of the image plane, most stereo methods produce semi-dense depth maps representing 3D scene edges. GTS produces modest results, albeit with many outliers. SGM has the most dense results because its regularizer fills in depth estimates in regions where the data fidelity term (time-image consistency) is not dominant. ESVO gives remarkable results in terms of accuracy and completeness, thus showing the effectiveness of its probabilistic inverse depth filters. Finally, our methods produce the best results: visually similar to ESVO but with finer details, thus being able to resolve more and finer edges in the scene.

Table 3 summarizes the quantitative performance with the metrics defined in Section 4.1 and on the same MVSEC sequences as [15], [38], [56], [68] (with ground truth depth). The best result per column is highlighted in bold, and the second best is underlined. Detailed, per-sequence tables are provided in the Supplementary Material. Contrary to previous works, we test on the entire sequences, consisting of 200 s (110 M events). The top part of Table 3 reports the results of the baseline methods, where ESVO is a top performer. EMVS is slightly worse than ESVO, consistently in most metrics, which shows that sensible stereo depth estimation (ESVO's Student-t filters) is beneficial to gain accuracy and reduce the number of outliers with respect to the monocular case. The bottom part of Table 3 reveals the results of several variations of Alg. 2 ((15)-(16)) with $N_s = 2$ sub-intervals. All accuracy and outlier metrics of Alg. 1 are significantly better than those of ESVO and EMVS, demonstrating the effectiveness of our fusion approaches: outperforming the state of the art and quantifying the gap between monocular and stereo methods.

Regarding completion, Alg. 1 recovers fewer points than ESVO. This has a natural explanation: ESVO generates several depth estimates per second that are propagated and fused. Even using ground truth poses, the estimates are noisy, and so they transfer to the fused image plane producing thick edges (Figure 8). By contrast, our stereo

method generates depth estimates at the AGT thresholding step, which is called just once, and therefore generates thinner edges than ESVO (finer details and well distributed at scene edges, as shown in Figure 8). To justify that lower completion values are not an issue, we applied a 4-neighbor morphological filter (MF) to dilate the mask of the depth map produced by AGT, and filled in the depth values using the center pixels. This almost quadrupled the number of reconstructed points (from 0.81M to 3M) while had a minimal effect on accuracy (row "MF" in Table 3), and therefore reduces the importance of comparing completion values given by very different algorithms. This also aligns with ideas in semi-dense and sparse SLAM, where fewer but more accurate points are preferred for several reasons: to have better distributed points on the image plane and to reduce the computational load (i.e., increase efficiency and speed) [69].

Variants $A_t \circ H_c$ and $A_c \circ A_t$ are also top performing. Despite the abundance of events accumulated in the DSI, Table 3 quantifies a gap between Alg. 1 and $A_c \circ A_t$ (the gap between the arithmetic mean and more concave fusion functions like *H* would be most noticeable if less data was used, as in Table 2). The last row of Table 3 is discussed in Section 4.4.3.

Driving sequences. The tests on the outdoor MVSEC sequences did not give good results because the camera baseline is very small compared to the scene depth, so geometrically the data is poor for 3D reconstruction. As noticed in [56], most points in those sequences are beyond the depth resolved by a disparity of 1 pix. Instead, we show results of our method on driving sequences from the DSEC dataset (Section 4.5), which has a larger baseline of 60 cm.

4.4 Temporal Fusion Experiments

4.4.1 Stereo

Figure 8 and Table 4 show the effect of temporal fusion. Using as few as $N_s = 2$ sub-intervals in Alg. 2 already delivers gains compared to the state of the art. While the optimal choice of the number of subintervals N_s depends

N_s	Mean Err [cm]↓	Median Err [cm]↓	bad-pix [%]↓	$\begin{array}{c} \text{SILog Err} \\ \times 100 \downarrow \end{array}$	AErrR [%]↓	$\begin{array}{c} \log \mathrm{RMSE} \\ \times 100 \downarrow \end{array}$	$\begin{array}{c} \delta < 1.25 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^2 \\ \llbracket \% \rrbracket \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^3 \\ [\%] \uparrow \end{array}$	#Points [million]↑
2	20.92	9.76	1.66	1.86	8.05	13.81	94.39	97.80	99.14	1.14
4	20.49	9.58	1.63	1.94	7.88	13.65	94.64	97.82	99.15	1.19
8	19.80	9.27	1.56	1.84	7.63	13.31	94.99	97.90	99.19	1.20

TABLE 4: Sensitivity of $A_t \circ H_c$ (Alg. 2) with respect to the number of subintervals N_s . Continuation of Table 3.



Fig. 9: Effect of temporal fusion ($N_s = 4$ sub-intervals) on monocular 3D reconstruction. Rocks scene in Figure 4.

on many factors, such as the number of events processed (duration of the intervals), the camera motion, etc. we found that using $N_s \in [2, 8]$ gives satisfactory results. This is reported in Table 4, which is an ablation study of Alg. 2 $A_t \circ H_c$ with respect to N_s . There is a clear trend: accuracy and completion values increase with N_s . However, memory and complexity also increases (linearly with N_s , see (19)). Comparing the values in Table 3 and Table 4, we notice that the improvement due to temporal fusion is not as pronounced as that due to stereo parallax (for example, the median error improves 32.00 % from EMVS to stereo $A_t \circ H_c$ ($N_s = 2$), and 5.02 % from $A_t \circ H_c$ $N_s = 2$ to $N_s = 8$).

4.4.2 Monocular

As a by-product, temporal fusion using the *H*-mean is a simple modification that can be applied to lightly improve the monocular method [10], specially when little data is available. While this is not the focus of the paper, we provide an example: Figure 9 shows the effect of monocular temporal fusion on the same slider rock-plane sequence as Figure 4. It compares the evolution of DSIs with and without temporal fusion. As expected, with *H*-mean fusion the DSI converges faster to the 3D structure and deemphasizes the locations of spurious ray intersections visible in the unfused DSI.

The effect of temporal fusion is dramatic if we compare the depth maps obtained from each subinterval with the depth map obtained after fusion. This is illustrated in Figure 10, where an interval of the *upenn_flying3* sequence [57] is divided into $N_s = 4$ sub-intervals. The first four columns depict depth maps produced by EMVS (row 1) and Alg. 1 (row 2) applied to individual sub-intervals (each with \approx 56k events per camera). These depth maps are noisy; however, when the DSIs are temporally fused and depth is extracted (Alg. 2), the final depth maps are remarkably cleaner (last column). Also, the fused stereo depth map has fewer outliers than the monocular one.

4.4.3 Stereo Fusion from Shuffled Sub-intervals

Arranging events in time subintervals allows us to question the event simultaneity assumption for depth estimation. Figure 11 shows the results of applying the system in Figure 3 with the shuffling block enabled to the MVSEC and UZH sequences. The shuffling block modifies line 4 in Alg 2 to use *different* subintervals for the S_1 fusion. The results are surprising: despite using non-corresponding subintervals (i.e., non-simultaneous events) for DSI fusion across cameras, the obtained depth- and confidence maps are very similar to those obtained with corresponding subintervals, with some added noise. Hence, event simultaneity is not needed for stereo depth estimation with our system (Figure 3). Only the similarity between the DSIs to be fused (intermediate ray density representations built by combining events and camera poses) is required.

Quantitatively, the last row of Table 3 informs about the depth errors and outliers incurred by shuffling. The differences with respect to the unshuffled case are small (slightly higher errors and outliers), which is remarkable given the quite diverse input events.

4.5 Experiments on DSEC Driving Dataset

We also give results on sequences from the driving dataset DSEC [16]. Figure 12 shows qualitative results. Driving scenarios are challenging for event-based sensors because forward motions typically produce considerably fewer events in the center of the image (where apparent motion is small) than in the periphery. Forward motion is also not particularly amenable for 3D reconstruction methods compared to sideways motions. Nevertheless, our stereo method shows notable results on this dataset. As expected, more 3D points are recovered in the periphery than in the center of the image, except when the car is turning.

Quantitative results are summarized in Table 5. Because the amount of events recorded by the VGA-resolution Prophesee Gen3 cameras is exorbitant, the experiments are carried out on a subset of the dataset. We test Alg. 1 and the baselines on 635 million events. Similarly to the results on MVSEC, Table 5 indicates that Alg. 1 has higher accuracy than ESVO and EMVS (at least 18.65% better in mean absolute error, and 42.30% in median absolute error). ESVO has marginally higher inlier values (1.45% for $\delta < 1.25$). The monocular method produces the largest errors (the mean and median absolute errors are 1.7-2.8× larger than those of stereo Alg. 1), and also a larger number of bad pixels and outliers. We also apply a morphological dilation filter (MF), with similar conclusions as in Table 3: the number of reconstructed points triples while the accuracy remains



Fig. 10: Effect of temporal fusion on the obtained depth maps. Depth is color coded as in Figure 8.



Fig. 11: Event Simultaneity? Fusion across cameras and time with shuffled time intervals ($N_s = 2$). Same scenes as Figure 8.

better than ESVO's. Finally, we notice that the ground truth depth of DSEC is sparser than that of MVSEC. This is due to the increased pixel resolution (VGA size) and the fact that LiDAR points do not fill as many camera pixels (percentage-wise) as in lower resolution cameras. Please see the accompanying video for a visual comparison between our method, ESVO, EMVS and GT.

4.6 Experiments on TUM-VIE Dataset

We present depth estimation results using Alg. 1 on the TUM-VIE dataset [59], the first public visual-inertial dataset with 1 Megapixel stereo event cameras (Prophesee Gen4 [5]). To the best of our knowledge, our work is the first to provide results on this new event-based dataset (the original paper presented the data but did not evaluate it on any event-based algorithm). Our experiments have served as a means to debug and fix the dataset. Since the dataset has no ground truth depth, we only present qualitative results.

Figure 13 presents results in indoor and outdoor sequences. Indoor sequences recorded in a room have ground truth poses given by a motion capture (mocap) system (columns 1 and 2). The indoor scene depth is small relative to the camera baseline (11.84 cm). Nevertheless, our method does a notable job in recovering 3D structure, with reduced number of outliers and clean the depth maps. Having set the reference view of the fused DSI on one camera trajectory, the large baseline makes the event rays back-projected from the other camera appear nearly parallel, which does not favor fusion.

For sequences recorded outside the mocap room, we computed ground truth poses using Basalt's VIO [70] on the stereo frames and IMU. The last two columns of Figure 13 depict the performance of our stereo method on such sequences. The space-sweeping method to build the DSIs works best with sideways translations that produce parallax necessary for the convergence of the back-projected rays. However, majority of the non-mocap sequences comprise forward camera motions, which contribute little parallax and also produce fewer events; hence they are not amenable for 3D reconstruction. The forward motion and the lower quality of camera poses lead to an overall poorer reconstruction quality compared to the sequences in the mocap room. On the other hand, the stereo setup is able to exploit the camera baseline as additional parallax for 3D reconstruction. In the accompanying video we provide a visual comparison between our method and ESVO on the TUM-VIE dataset, showing that our method performs better at higher resolutions than ESVO.

We noticed that the *calibrationA* sequences (*skate-easy*, *desk2*) produced better results than the *calibrationB* sequences, which leads us to believe that the calibration errors



Fig. 12: *Results on DSEC data* [16]. Top row: Semi-dense depth maps (overlaid on color frames) estimated by Alg. 1 on event packets of 200 ms. Depth is color-coded from red (close) to blue (far), in the range 4–200 m. Bottom row: confidence maps.

Algorithm	Mean Err [m]↓	Median Err [m]↓	bad-pix [%]↓	$\begin{array}{c} \text{SILog Err} \\ \times 100 \downarrow \end{array}$	AErrR [%]↓	$\begin{array}{c} \log \mathrm{RMSE} \\ \times 100 \downarrow \end{array}$	$\begin{array}{c} \delta < 1.25 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^2 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^3 \\ [\%] \uparrow \end{array}$	#Points [million]↑
EMVS [10] (monocular)	5.64	2.52	13.68	13.23	25.52	36.49	72.56	87.12	93.56	1.31
ESVO [15]	3.88	1.56	12.08	9.23	18.89	30.80	84.53	92.57	95.63	3.40
$H_c \circ A_t \text{ (Alg 1)}$	3.27	0.90	10.75	8.19	17.48	28.73	<u>83.30</u>	<u>91.56</u>	<u>95.62</u>	1.25
$H_c \circ A_t (\text{Alg 1}) + \text{MF}$	3.51	<u>0.96</u>	<u>11.81</u>	8.89	18.84	29.99	81.72	90.68	95.07	3.83

TABLE 5: Quantitative evaluation on the driving dataset DSEC (zurich04a sequence) with maximum ground truth depth 50 m. The methods are evaluated on 35s of stereo data, consisting of 635 million events and containing 350 ground truth depth maps. Each depth map is computed using 0.2 s of event data (≈ 3.5 million events). ESVO is executed fusing two depth maps generated at 10 Hz (LiDAR rate), i.e., 0.2 s of event data. MF: morphological filter.

were significant in the latter. We also compared the results of our method with two different sources of ground truth poses (the mocap system and Basalt), and observed no significant differences in the depth- and confidence maps. Hence, we concluded that (*i*) the poses from Basalt may be considered as accurate as the mocap for short time intervals (e.g., 0.5 s, with \approx 10M events), and (*ii*) there is robustness to noise: our methods provide reasonable depth estimates using poses from a VIO (non-mocap) algorithm.

Event cameras offer advantages over frame-based cameras to handle HDR scenes, as demonstrated in Figure 14. In the first row, the grayscale frame is overexposed in the outdoor area, whereas the events capture the floor tiles and garden scene well. In the second row, the end of the corridor is underexposed in the frame, whereas the events capture the whole scene. Our stereo method correctly estimates depth in all regions due to the HDR capabilities of the events.

4.7 Effect of Varying the Sensor's Spatial Resolution

So far, results on datasets from three different resolutions have been presented: DAVIS346 (MVSEC), Prophesee Gen3 (DSEC) and Prophesee Gen4 (TUM-VIE). However, the scenes are all different, some do not have ground truth depth, and calibration errors may influence the results. To analyze the response of our stereo Alg. 1 to varying pixel resolutions under controlled conditions (same scene, same FOV, etc.), we generate events using a simulator (ESIM [60]). We use the textured scene *flying_room*, with a camera baseline of 20 cm and the OpenGL renderer. The scene is visualized in Figure 15 (bottom right). Inspired by MVSEC,

DSEC and TUM-VIE datasets, three sensor resolutions are tested: 320×240 pix, 640×480 pix and 1280×960 pix, respectively.

We run Alg. 1 on stereo events generated during the same time duration (increasing the sensor resolution increases the number of events generated within the same time duration: roughly 0.25 Mev, 2.25 Mev and 14.3 Mev for the three above resolutions, respectively). Figure 15 compares qualitatively the semi-dense depth maps and confidence maps recovered by our method on the three sensor resolutions inputs. The observations from these images are complemented by the quantitative analysis in Figure 16. Figure 16a demonstrates quantitatively that depth errors decrease as the sensor resolution increases, for both monocular and stereo cases. Stereo fusion reduces the error almost by half with respect to the monocular case. Next, we also analyze the density of the reconstructions. Figure 16 also reports the precision, recall and F1-score plots for the depth maps in Figure 15. A clear trend is observed: the precision increases with increasing camera resolution (lower errors, like in Figure 16a). For the highest resolution tested $(1280 \times 960 \text{ pix})$, the precision reached $\approx 97 \%$ within 4 cm of depth error tolerance. However, the recall curves state that a lower resolution may allow us to recover more 3D points in the scene (6.00 % vs. 2.00 %), at the expense of bigger depth errors. Since the depth outputs are highly sparse due to the nature of event data, the F1-score is heavily skewed by the low values of recall.



Fig. 13: *Depth estimation using 1Mpix event cameras.* Depth estimated using Alg 1 on 0.50 s intervals of data from TUM-VIE [59]. Depth maps are color-coded from red (near) to blue (far). The range is 0.45–4 m for indoor sequences in the motion-capture room (*6dof, desk2*), 1–20 m for the corridor sequence *skate-easy* and 3–200 m for the outdoor sequence *bike-easy*.



Fig. 14: *HDR scenes.* Output of Alg. 1 on HDR scenes from the TUM-VIE dataset. Unlike frame-based cameras, event cameras can perceive both under- and over-exposed regions of the scene well, leading to good depth estimation throughout.

4.8 Fusing More than Two Event Cameras

We also test our method on sequences from the EVIMO2 dataset [71], recorded with a *trinocular* event-camera rig consisting of a Samsung DVS Gen3 [72] and two Prophesee CD Gen3 event cameras [73], all with 640×480 pix. The field of views (FOVs) of the cameras have a narrow overlap due to the way they are arranged in the sensor rig (Prophesee event cameras are in portrait mode, whereas the Samsung DVS, in the middle, is in landscape mode). We set the central camera (Samsung DVS) as the reference one.

Figure 17 shows the depth- and confidence maps from each camera separately and for the fused DSI during two motions (two pairs of columns): a normal one and a fast motion, with retinal speeds between 1900 pix/s for objects in the far end and 3500 pix/s for objects close to the camera. The resulting semi-dense depth map obtained from the fused DSI inherit the above-mentioned narrow FOV overlap of the cameras. Overall, the fused depth map suppresses noise that would otherwise appear as very prominent outliers in the individual depth maps. This experiment shows that *our method naturally fuses multiple cameras* with linear

TABLE 6: Runtime comparison of the different steps of three algorithms. Parameters: $N_e \approx 235$ kev, $N_c = 2$, $N_p \approx 90$ k (DAVIS346), $N_Z = 100$, $N_k = 25$ (5 × 5), $N_s = 2$ for Alg. 2 (here, $A_t \circ H_c$). Time is given in ms.

EMVS [10]	Alg. <mark>1</mark>	Alg. 2
233.96	520	0.01
-	51.09	98.01
36.34	35	.62
0.28	0.	28
	EMVS [10] 233.96 - 36.34 0.28	EMVS [10] Alg. 1 233.96 520 - 51.09 36.34 35 0.28 0.

complexity, i.e., without handling them by pairs, as prior works do.

4.9 Runtime

Complementing the complexity analysis in Section 3.7, Table 6 presents the average time taken in each step (DSI creation, DSI fusion, argmax and AGT thresholding) over 100 sample runs, on a laptop with an Intel i7-10510U 8-core CPU. We consider typical numbers from a stereo DAVIS346 configuration [74]. For comparison, we also report the num-



Fig. 15: *Effect of the sensor's spatial resolution*. Output of Alg. 1 on events simulated with different camera resolutions. The scene *–flying_room–* is shown on the top right. Depth maps are colored from red (near) to blue (far) in the range 1.3–4 m.



Fig. 16: (a) Depth errors across various camera resolutions for monocular [10] and stereo Alg. 1 methods. Maximum scene depth is 2.70 m. (b)-(d) Precision, recall and F1-score for depth estimated using Alg. 1 on scenes of different resolutions.

bers obtained with only one of the cameras (monocular setup [10]). DSI creation takes the longest time as it is a complex transformation and depends on the number of input events. We did not find major runtime differences in our implementation of the fusion functions tested (6)–(11).

The numbers in Table 6 agree with complexity formulas (17)–(19). The DSI creation runtime of the stereo methods is roughly twice ($N_c = 2$) that of the monocular method (520 vs. 234 ms). The temporal fusion step with $N_s = 2$ sub-intervals is twice as expensive as that with one interval (98 vs. 51 ms). The arg max and AGT steps are the same for all methods, thus no major runtime differences are observed.

4.10 Effect of Varying the Contrast Sensitivity

Due to the sparse nature of event data, which is largely controlled in the camera by the contrast sensitivity threshold, it is interesting to analyze the performance of our 3D reconstruction method for various sparsity levels. To this end, we ran Alg. 1 on stereo events simulated using ESIM [60] with five levels of event generation contrast threshold $\theta = \{0.05, 0.1, 0.2, 0.4, 0.8\}$. A small contrast threshold implies that a small change in brightness is sufficient to trigger events, and thus leads to the generation of many events from edges and textures in the scene (i.e., high sensor

sensitivity). Figure 18 illustrates the 3D reconstruction quality across varying contrast thresholds for the flying room sequence. In general, we observe stable reconstructions as the contrast threshold increases, except for the fact that fewer points are recovered (e.g., on the floor) and more noisy points (i.e., outliers) appear. This is also observed quantitatively in Figure 19a, where the mean error increases for high contrast thresholds (the larger number of outliers distort the mean error) while the median error remains fairly constant (a sign of "stability" if outliers are removed). Figures 19b–19d depict the precision, recall and F1-score curves for depth predicted using data from different contrast thresholds. We observe that low contrast thresholds provide better precision and recall since they have fewer noisy outliers and recover more 3D points respectively. The increase in precision as well as recall comes at the cost of increased computational overhead as more events need to be processed at lower contrast thresholds.

In summary, the synthetic experiments suggest that increasing the event count either by decreasing the contrast threshold or increasing the spatial resolution of the camera improves 3D reconstruction accuracy. This comes at the cost of increased computational effort needed to process more input data.



(a) Trinocular rig from EVIMO2v1. Used with permission [71], 2019.



Fig. 17: *Trinocular event-camera fusion*. Output of Stereo Alg. 1 on the SFM sequence 03_00 from the EVIMO2 dataset [71]. The first row shows an RGB image of the scene and the events (from the Samsung DVS [72]) over a short time duration, displayed in red/blue according to polarity. The following rows show the output of our stereo method using the three event cameras in the dataset.

While event cameras are dominated by those that compute temporal contrast [4] (DVS), an "event" could have a broader interpretation, such as any meaningful information that decreases demands on bandwidth, memory, and power for data transmission, storage and processing. This work has analyzed the advantages that event cameras offer for stereo depth estimation. Varying more than the temporal contrast is possible with different prototype vision sensors, such as Parallel Processor Arrays (PPAs). These sensors embed a processor within each pixel and are thus programmable, enabling a richer family of operations (at the expense of larger pixel sizes or more transistor layers) [75].

4.11 Discussion

Let us summarize some of our findings. In the UZH dataset (Figure 8, top three rows), our method achieves best results compared to the state of the art (SOTA). Figure 8 also shows that our stereo method, in its different variations, recovers depth at more fine-detailed structures than SOTA method ESVO.

In the indoor MVSEC dataset, Table 3 shows that our stereo method outperforms SOTA quantitatively across multiple standard metrics. We also show that the effect of time fusion, while being significant because it speeds up structure convergence and cleans up depth maps, has a smaller effect than that of fusion induced by parallax from an additional sensor (i.e., switching from monocular to stereo) (Tables 3 and 4). Remarkably, both strategies are unified in the same theory of fusion of refocused events that we propose.

The outdoor driving MVSEC sequences have very small stereo baseline, which makes them poor for 3D reconstruction [76], hence we test on the recent DSEC dataset. Here, our method also outperforms the SOTA method ESVO (Table 5).

The TUM-VIE dataset allowed us to demonstrate experiments on high event camera resolution (1Mpix) and robustness to errors in the camera poses. The EVIMO2 dataset allowed us to establish multi-camera (trinocular) depth estimation and during high-speed motion (which blurs regular frames).

Throughout the experiments (MVSEC, DSEC, TUM-VIE, time fusion, etc.), we have shown the gains with respect to the monocular method; the main advantages of stereo are: higher accuracy, outlier rejection, and faster convergence (due to the additional parallax).

Additionally, we have analyzed the sensitivity with respect to the camera's spatial resolution and contrast threshold: the higher the resolution or the lower the threshold, the more accurate our method becomes, at the expense of computational burden due to the larger number of input events. We also analyzed the computational performance, showing the agreement between complexity (theory) and runtime (practice).

Most interestingly, we have analyzed the effect of shuffling events: our method does not need event simultaneity. It can fuse DSIs even if they are built from events well separated in time. The best results are obtained fusing identical intervals.

4.12 Limitations

Our method assumes, like multi-view 3D reconstruction methods for standard cameras, accurate calibration and camera poses. These assumptions allowed us to concentrate our efforts on investigating fusion functions for eventbased 3D reconstruction under good multi-view alignment conditions. Calibration is today of good quality using the DAVIS frames or image reconstruction [62] if frames are not available. In a full SLAM system, noise in the poses can propagate to the 3D reconstruction module. The TUM-VIE



Fig. 18: *Effect of varying the camera's contrast threshold*. Output of stereo Alg. 1 on simulated events generated with varying contrast threshold θ . As θ increases, the sparsity of recovered points increases and the quality of depth errors decreases. Depth maps are colored from red (near) to blue (far) in the range 1.3–4 m.



Fig. 19: (a) Sensitivity of depth errors with respect to the sensor's contrast threshold. Maximum scene depth is 2.70 m. (b)-(d) Precision, recall and F1-score for depth estimated using Alg. 1 using events simulated with different contrast thresholds.

experiments, with different pose sources (mocap vs. VIO algorithm), showed the robustness of our method to realistic poses (i.e., noisy as produced by a VIO algorithm). This is encouraging, as future research on event-based camera tracking could achieve such accuracy and be combined with our stereo method. Recent results on the monocular case [77] suggests that such accuracy and robustness is achievable combining events and frames.

Another assumption of our method is that events are dominantly triggered by moving edges (brightness constancy). Hence, it may fail to estimate depth accurately from events that are not due to motion, such as those caused by flickering lights. Such events may be removed during preprocessing [78].

5 CONCLUSION

We presented simple and effective state-of-the-art algorithms for event-based multi-camera 3D reconstruction in SLAM, combining across-camera and temporal fusion. Thanks to the availability of accurate camera poses, matching within and across event cameras happens *implicitly* in DSI space, which removes the need for *event simultaneity* (explicit data association). We investigated DSI space fusion methods, and showed how *the same technique unifies temporal and camera fusion*. We proposed a spectrum of fusion functions, including summing event rays in multi-camera settings, and objectively analyzed the results they produced. Our theoretical design was supported by a comprehensive set of experiments: testing on five established datasets and a simulator on a variety of scenarios, on millions of input events, while outperforming state of the art methods. Fusion functions like the *H*-mean are beneficial for fusion because they have strong concavity, are bounded and smooth. The effect of the parallax given by an additional camera is stronger than the effect of temporal fusion. Our method works well regardless of the spatial resolution, which is interesting given the increasingly high spatial resolution of event cameras (see the comparisons with ESVO on DSEC and TUM-VIE datasets in the accompanying video).

Future research may look into combining the proposed method with an appropriately designed camera tracking algorithm to yield a full event-based stereo visual SLAM pipeline. The results on the driving datasets open the door to applying the proposed technique for creating HDR edge maps of a vehicle's surroundings and using it for later processing stages of Spatial AI and mobile intelligent systems [79], such as spatial awareness or extraction of semantic information.

APPENDIX

Tables 7 to 9 report depth estimation metrics on individual sequences of the MVSEC dataset. Table 3 is the average of these per-sequence tables.

	Algorithm	Mean Err [cm]↓	Median Err [cm]↓	bad-pix [%]↓	$\begin{array}{c} \text{SILog Err} \\ \times 100 \downarrow \end{array}$	AErrR [%]↓	$\begin{array}{c} \log \mathrm{RMSE} \\ \times 100 \downarrow \end{array}$	$\begin{array}{c} \delta < 1.25 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^2 \\ \llbracket \% \rrbracket \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^3 \\ [\%] \uparrow \end{array}$	#Points [million]↑
SOTA	EMVS [10] (monocular) ESVO [15] ESVO indep. 1s SGM indep. 1s GTS indep. 1s	39.37 24.04 23.39 35.45 700.37	14.95 10.21 10.03 13.61 38.39	3.05 2.54 2.18 5.54 32.51	4.72 2.94 2.79 7.35 79.26	13.25 9.76 9.78 15.03 111.21	22.10 17.17 16.72 27.46 91.44	82.03 91.43 91.57 85.96 54.27	93.43 96.53 96.84 93.51 67.16	97.62 98.55 98.79 96.40 73.39	$1.21 \\ 1.95 \\ 1.41 \\ 11.64 \\ 0.03$
Ours	$ \begin{array}{l} H_c \circ A_t \ (\text{Alg 1}) \\ H_c \circ A_t \ (\text{Alg 1}) + \text{MF} \\ H_c \circ H_t \\ H_t \circ A_c \\ A_c \circ H_t \\ A_c \circ A_t \\ A_t \circ H_c \\ A_t \circ H_c \\ A_t \circ H_c + \text{shuffling} \end{array} $	22.53 23.33 24.38 23.78 24.48 21.93 22.32 23.92	$\begin{array}{c} 9.72\\ 9.90\\ 10.81\\ 10.37\\ 10.78\\ 9.35\\ 9.63\\ 10.60\\ \end{array}$	1.30 1.39 1.40 1.26 1.39 1.22 1.31 1.28	1.94 2.08 2.11 2.03 2.14 1.87 1.92 2.01	7.91 8.12 8.59 8.36 8.62 7.75 7.85 8.40	$14.11 \\ 14.61 \\ 14.71 \\ 14.41 \\ 14.81 \\ 13.85 \\ 14.02 \\ 14.36$	93.49 93.16 92.53 92.72 92.48 93.65 93.62 92.76	97.50 97.28 97.38 97.54 97.34 97.65 97.51 97.55	99.17 99.05 99.15 99.21 99.13 99.23 99.17 99.21	$\begin{array}{c} 0.96\\ 3.48\\ 1.24\\ 1.02\\ 1.20\\ 0.87\\ 1.03\\ 1.11\\ \end{array}$

TABLE 7: Quantitative comparison of the proposed methods with the state of the art. MVSEC flying1. See Table 3.

	Algorithm	Mean Err [cm]↓	Median Err [cm]↓	bad-pix [%]↓	$\begin{array}{c} \text{SILog Err} \\ \times 100 \downarrow \end{array}$	AErrR [%]↓	$\begin{array}{c} \log \mathrm{RMSE} \\ \times 100 \downarrow \end{array}$	$\begin{array}{c} \delta < 1.25 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^2 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^3 \\ [\%] \uparrow \end{array}$	#Points [million]↑
SOTA	EMVS [10] (monocular) ESVO [15] ESVO indep. 1s SGM indep. 1s GTS indep. 1s	31.42 21.34 20.42 32.94 167.14	13.01 8.97 8.63 8.75 37.23	6.15 3.75 3.50 8.29 43.08	4.56 3.48 3.24 9.50 71.91	13.37 9.32 9.14 15.82 94.78	21.80 19.14 18.35 31.54 86.93	84.07 91.60 92.03 84.40 49.36	94.72 95.88 96.19 92.33 60.54	97.88 97.86 98.19 95.48 67.76	$1.17 \\ 1.89 \\ 1.41 \\ 16.95 \\ 0.07$
Ours	$ \begin{array}{l} H_c \circ A_t \ (\text{Alg 1}) \\ H_c \circ A_t \ (\text{Alg 1}) + \text{MF} \\ H_c \circ H_t \\ H_t \circ A_c \\ A_c \circ H_t \\ A_c \circ A_t \\ A_t \circ H_c \\ A_t \circ H_c + \text{shuffling} \end{array} $	18.20 18.58 22.47 21.37 22.28 19.04 19.82 21.66	8.49 8.68 10.13 9.86 10.06 8.76 8.91 9.91	1.77 1.86 2.92 2.50 2.81 2.24 2.55 2.60	$ 1.78 \\ 1.81 \\ 2.46 \\ 2.26 \\ 2.40 \\ 1.99 \\ 2.09 \\ 2.22 $	8.13 8.19 9.41 9.13 9.35 8.44 8.60 9.16	13.59 13.71 16.07 15.41 15.88 14.40 14.77 15.27	95.53 95.27 92.66 93.41 92.80 94.68 94.16 93.14	98.13 98.07 97.19 97.60 97.28 97.89 97.55 97.53	99.08 99.09 98.79 98.92 98.82 99.02 98.94 98.94 98.92	$\begin{array}{c} 0.65 \\ 2.42 \\ 1.39 \\ 1.18 \\ 1.34 \\ 1.01 \\ 1.17 \\ 1.22 \end{array}$

TABLE 8: Quantitative comparison of the proposed methods with the state of the art. MVSEC flying2. See Table 3.

	Algorithm	Mean Err [cm]↓	Median Err [cm]↓	bad-pix [%]↓	SILog Err ×100↓	AErrR [%]↓	$\begin{array}{c} \log \mathrm{RMSE} \\ \times 100 \downarrow \end{array}$	$\begin{array}{c} \delta < 1.25 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^2 \\ [\%] \uparrow \end{array}$	$\begin{array}{c} \delta < 1.25^3 \\ [\%] \uparrow \end{array}$	#Points [million]↑
SOTA	EMVS [10] (monocular) ESVO [15] ESVO indep. 1s SGM indep. 1s GTS indep. 1s	30.54 29.62 24.29 37.86 299.48	15.09 12.61 10.84 14.69 60.66	2.31 3.78 2.81 5.33 39.75	3.33 4.02 3.05 8.52 72.24	11.59 11.50 9.84 17.65 102.77	18.27 20.20 17.54 29.46 88.87	88.16 88.28 91.87 85.67 45.04	96.45 94.88 96.46 93.31 58.86	98.47 97.52 98.16 96.21 66.94	1.42 2.29 1.86 14.81 0.08
Ours	$ \begin{array}{l} H_c \circ A_t \ (\text{Alg 1}) \\ H_c \circ A_t \ (\text{Alg 1}) + \text{MF} \\ H_c \circ H_t \\ H_t \circ A_c \\ A_c \circ H_t \\ A_c \circ A_t \\ A_t \circ H_c \\ A_t \circ H_c + \text{shuffling} \end{array} $	19.49 20.02 23.48 22.68 23.29 20.17 20.61 22.24	10.38 10.59 11.99 11.83 11.97 10.68 10.74 11.60	$\begin{array}{c} 0.99\\ 1.02\\ 1.36\\ 1.24\\ 1.32\\ 1.06\\ 1.13\\ 1.17\end{array}$	$1.43 \\ 1.50 \\ 1.95 \\ 1.81 \\ 1.91 \\ 1.53 \\ 1.59 \\ 1.73$	7.35 7.50 8.58 8.35 8.52 7.59 7.68 8.17	12.01 12.30 14.01 13.50 13.86 12.40 12.64 13.20	96.09 95.79 93.79 94.33 93.94 95.67 95.40 94.58	98.60 98.51 97.90 98.14 97.97 98.48 98.35 98.21	99.38 99.36 99.14 99.24 99.17 99.36 99.30 99.27	0.82 3.11 1.79 1.55 1.72 1.09 1.23 1.39

TABLE 9: Quantitative comparison of the proposed methods with the state of the art. MVSEC flying3. See Table 3.

REFERENCES

- D. Berco and D. Shenp Ang, "Bioinspired robotic vision with online learning capability and rotation-invariant properties," *Adv. Intell. Syst.*, vol. 3, no. 8, p. 2100025, 2021.
- [2] G. Gallego, T. Delbruck, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza, "Event-based vision: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 1, pp. 154–180, 2022.
- [3] D. Berco and D. Shenp Ang, "Recent progress in synaptic devices paving the way toward an artificial cogni-retina for bionic and machine vision," *Adv. Intell. Syst.*, vol. 1, no. 1, p. 1900003, 2019.
- [4] P. Lichtsteiner, C. Posch, and T. Delbruck, "A 128×128 120 dB 15

 μ s latency asynchronous temporal contrast vision sensor," *IEEE J.* Solid-State Circuits, vol. 43, no. 2, pp. 566–576, 2008.

- [5] T. Finateu, A. Niwa, D. Matolin, K. Tsuchimoto, A. Mascheroni, E. Reynaud, P. Mostafalu, F. Brady, L. Chotard, F. LeGoff, H. Takahashi, H. Wakabayashi, Y. Oike, and C. Posch, "A 1280x720 backilluminated stacked temporal contrast event-based vision sensor with 4.86µm pixels, 1.066Geps readout, programmable event-rate controller and compressive data-formatting pipeline," in *IEEE Intl. Solid-State Circuits Conf. (ISSCC)*, 2020, pp. 112–114.
- [6] Y. Suh, S. Choi, M. Ito, J. Kim, Y. Lee, J. Seo, H. Jung, D.-H. Yeo, S. Namgung, J. Bong, J. seok Kim, P. K. J. Park, J. Kim, H. Ryu, and Y. Park, "A 1280x960 Dynamic Vision Sensor with a 4.95-μm pixel pitch and motion artifact minimization," in *IEEE Int. Symp. Circuits Syst. (ISCAS)*, 2020, pp. 1–5.

- [7] T. Delbruck and M. Lang, "Robotic goalie with 3ms reaction time at 4% CPU load using event-based dynamic vision sensor," Front. Neurosci., vol. 7, p. 223, 2013.
- N. Matsuda, O. Cossairt, and M. Gupta, "MC3D: Motion contrast [8] 3D scanning," in IEEE Int. Conf. Comput. Photography (ICCP), 2015, pp. 1-10.
- [9] A. Amir, B. Taba, D. Berg, T. Melano, J. McKinstry, C. D. Nolfo, T. Nayak, A. Andreopoulos, G. Garreau, M. Mendoza, J. Kusnitz, M. Debole, S. Esser, T. Delbruck, M. Flickner, and D. Modha, "A low power, fully event-based gesture recognition system," in IEEE Conf. Comput. Vis. Pattern Recog. (CVPR), 2017, pp. 7388-7397.
- [10] H. Rebecq, G. Gallego, E. Mueggler, and D. Scaramuzza, "EMVS: Event-based multi-view stereo-3D reconstruction with an event camera in real-time," Int. J. Comput. Vis., vol. 126, no. 12, pp. 1394-1414, Dec. 2018.
- [11] W. Wang, E. Covi, A. Milozzi, M. Farronato, S. Ricci, C. Sbandati, G. Pedretti, and D. Ielmini, "Neuromorphic motion detection and orientation selectivity by volatile resistive switching memories," Adv. Intell. Syst., vol. 3, no. 4, p. 2000224, 2021.
- [12] M. E. Sayed, J. O. Roberts, K. Donaldson, S. T. Mahon, F. Iqbal, B. Li, S. Franco Aixela, G. Mastorakis, E. T. Jonasson, M. P. Nemitz, S. Bernardini, and A. A. Stokes, "Modular robots for enabling operations in unstructured extreme environments," Adv. Intell. Syst., vol. 4, no. 5, p. 2000227, 2022.
- [13] H. Rebecq, T. Horstschäfer, G. Gallego, and D. Scaramuzza, "EVO: A geometric approach to event-based 6-DOF parallel tracking and mapping in real-time," IEEE Robot. Autom. Lett., vol. 2, no. 2, pp. 593-600, 2017.
- [14] H. Kim, S. Leutenegger, and A. J. Davison, "Real-time 3D reconstruction and 6-DoF tracking with an event camera," in Eur. Conf. Comput. Vis. (ECCV), 2016, pp. 349–364. [15] Y. Zhou, G. Gallego, and S. Shen, "Event-based stereo visual
- odometry," IEEE Trans. Robot., vol. 37, no. 5, pp. 1433-1450, 2021.
- [16] M. Gehrig, W. Aarents, D. Gehrig, and D. Scaramuzza, "DSEC: A stereo event camera dataset for driving scenarios," IEEE Robot. Autom. Lett., vol. 6, no. 3, pp. 4947–4954, 2021.
- [17] G. Gallego, H. Rebecq, and D. Scaramuzza, "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), 2018, pp. 3867-3876.
- [18] G. Gallego, M. Gehrig, and D. Scaramuzza, "Focus is all you need: Loss functions for event-based vision," in IEEE Conf. Comput. Vis. Pattern Recog. (CVPR), 2019, pp. 12272–12281.
- [19] T. Stoffregen and L. Kleeman, "Event cameras, contrast maximization and reward functions: an analysis," in IEEE Conf. Comput. Vis. Pattern Recog. (CVPR), 2019, pp. 12292-12300.
- [20] T. Stoffregen, G. Gallego, T. Drummond, L. Kleeman, and D. Scaramuzza, "Event-based motion segmentation by motion compensation," in Int. Conf. Comput. Vis. (ICCV), 2019, pp. 7243-7252
- [21] Z. W. Wang, P. Duan, O. Cossairt, A. Katsaggelos, T. Huang, and B. Shi, "Joint filtering of intensity images and neuromorphic events for high-resolution noise-robust imaging," in IEEE Conf. Comput. Vis. Pattern Recog. (CVPR), 2020, pp. 1606-1616.
- [22] U. M. Nunes and Y. Demiris, "Robust event-based vision model estimation by dispersion minimisation," IEEE Trans. Pattern Anal. Mach. Intell., 2021.
- [23] X. Peng, L. Gao, Y. Wang, and L. Kneip, "Globally-optimal contrast maximisation for event cameras," IEEE Trans. Pattern Anal. Mach. Intell., vol. 44, no. 7, pp. 3479-3495, 2022.
- [24] S. Shiba, Y. Aoki, and G. Gallego, "Secrets of event-based optical flow," in Eur. Conf. Comput. Vis. (ECCV), 2022.
- [25] J. J. Hagenaars, F. Paredes-Valles, and G. C. H. E. de Croon, "Self-supervised learning of event-based optical flow with spiking neural networks," in Advances in Neural Information Processing Systems (NeurIPS), vol. 34, 2021, pp. 7167–7179.
- [26] S. Shiba, Y. Aoki, and G. Gallego, "Event collapse in contrast maximization frameworks," Sensors, vol. 22, no. 14, pp. 1-20, 2022.
- [27] H. Seok and J. Lim, "Robust feature tracking in dvs event stream using Bezier mapping," in IEEE Winter Conf. Appl. Comput. Vis. (WACV), 2020, pp. 1647–1656.
- [28] S. Gao, Y. Dai, and A. Nathan, "Tactile and Vision Perception for Intelligent Humanoids," Adv. Intell. Syst., vol. 4, no. 2, p. 2100074, 2022
- [29] S. Ghosh, G. D'Angelo, A. Glover, M. Iacono, E. Niebur, and C. Bartolozzi, "Event-driven proto-object based saliency in 3D space to attract a robot's attention," Scientific Reports, vol. 12, no. 1, p. 7645, Dec 2022.

- [30] M. Mahowald, "VLSI analogs of neuronal visual processing: A synthesis of form and function," Ph.D. dissertation, California Institute of Technology, Pasadena, California, May 1992.
- [31] -, The Silicon Retina. Boston, MA: Springer US, 1994, pp. 4-65. [32] D. Marr and T. Poggio, "Cooperative computation of stereo disparity," Science, vol. 194, no. 4262, pp. 283-287, 1976.
- [33] E. Piatkowska, A. N. Belbachir, and M. Gelautz, "Asynchronous stereo vision for event-driven dynamic stereo sensor using an adaptive cooperative approach," in Int. Conf. Comput. Vis. Workshops (ICCVW), 2013, pp. 45-50.
- [34] M. Firouzi and J. Conradt, "Asynchronous event-based cooperative stereo matching using neuromorphic silicon retinas," Neural Proc. Lett., vol. 43, no. 2, pp. 311-326, 2016.
- [35] M. Osswald, S.-H. Ieng, R. Benosman, and G. Indiveri, "A spiking neural network model of 3D perception for event-based neuromorphic stereo vision systems," Sci. Rep., vol. 7, no. 1, Jan. 2017.
- [36] L. Steffen, D. Reichard, J. Weinland, J. Kaiser, A. Rönnau, and R. Dillmann, "Neuromorphic stereo vision: A survey of bioinspired sensors and algorithms," Front. Neurorobot., vol. 13, p. 28, 2019.
- J. Furmonas, J. Liobe, and V. Barzdenas, "Analytical review of [37] event-based camera depth estimation methods and systems," Sensors, vol. 22, no. 3, p. 1201, Feb. 2022
- [38] Y. Zhou, G. Gallego, H. Rebecq, L. Kneip, H. Li, and D. Scaramuzza, "Semi-dense 3D reconstruction with a stereo event camera," in Eur. Conf. Comput. Vis. (ECCV), 2018, pp. 242-258.
- [39] J. N. P. Martel, J. Müller, J. Conradt, and Y. Sandamirskaya, "An active approach to solving the stereo matching problem using event-based sensors," in IEEE Int. Symp. Circuits Syst. (ISCAS), 2018, pp. 1-5.
- [40] M. Muglikar, G. Gallego, and D. Scaramuzza, "ESL: Event-based structure light," in Int. Conf. 3D Vision (3DV), 2021, pp. 1165-1174.
- [41] G. Haessig, X. Berthelon, S.-H. Ieng, and R. Benosman, "A spiking neural network model of depth from defocus for event-based neuromorphic vision," Sci. Rep., vol. 9, no. 3744, 2019.
- J. Carneiro, S.-H. Ieng, C. Posch, and R. Benosman, "Event-based [42]3D reconstruction from neuromorphic retinas," Neural Netw., vol. 45, pp. 27–38, 2013.
- [43] S.-H. Ieng, J. Carneiro, M. Osswald, and R. Benosman, "Neuromorphic event-based generalized time-based stereovision," Front. Neurosci., vol. 12, p. 442, 2018.
- R. Benosman, S.-H. Ieng, C. Clercq, C. Bartolozzi, and M. Srini-[44]vasan, "Asynchronous frameless event-based optical flow," Neural Netw., vol. 27, pp. 32-37, 2012.
- [45] E. Piatkowska, A. N. Belbachir, and M. Gelautz, "Cooperative and asynchronous stereo vision for dynamic vision sensors," Meas. Sci. Technol., vol. 25, no. 5, p. 055108, Apr. 2014.
- [46] D. Gehrig, H. Rebecq, G. Gallego, and D. Scaramuzza, "EKLT: Asynchronous photometric feature tracking using events and frames," Int. J. Comput. Vis., vol. 128, pp. 601-618, 2020.
- [47] X. Lagorce, G. Orchard, F. Gallupi, B. E. Shi, and R. Benosman, "HOTS: A hierarchy of event-based time-surfaces for pattern recognition," IEEE Trans. Pattern Anal. Mach. Intell., vol. 39, no. 7, pp. 1346-1359, Jul. 2017.
- [48] R. T. Collins, "A space-sweep approach to true multi-image matching," in IEEE Conf. Comput. Vis. Pattern Recog. (CVPR), 1996, pp. 358-363.
- [49] D. Yuan and R. Chung, "Correspondence-free stereo vision for the case of arbitrarily-positioned cameras," in IEEE Int. Conf. Robot. Autom. (ICRA), vol. 2, 2003, pp. 1688–1693 vol.2.
- [50] S. Lehmann, A. P. Bradley, J. Williams, P. J. Kootsookos, and L. Clarkson, "Correspondence-free determination of the affine fundamental matrix," IEEE Trans. Pattern Anal. Mach. Intell., vol. 29, no. 1, pp. 82-97, 2007.
- [51] J. Aloimonos and J.-Y. Herve, "Correspondenceless stereo and motion: planar surfaces," IEEE Trans. Pattern Anal. Mach. Intell., vol. 12, no. 5, pp. 504–510, 1990.
- [52] H. Cho, J. Jeong, and K.-J. Yoon, "EOMVS: Event-based omnidirectional multi-view stereo," IEEE Robot. Autom. Lett., vol. 6, no. 4, pp. 6709-6716, Oct. 2021.
- [53] C. Brandli, R. Berner, M. Yang, S.-C. Liu, and T. Delbruck, "A 240x180 130dB 3µs latency global shutter spatiotemporal vision sensor," IEEE J. Solid-State Circuits, vol. 49, no. 10, pp. 2333-2341, 2014.
- [54] C.-A. Deledalle, L. Denis, and F. Tupin, "How to compare noisy patches? Patch similarity beyond Gaussian noise." Int. J. Comput. *Vis.*, vol. 99, no. 1, pp. 86–102, 2012.

- [55] G. Gallego and D. Scaramuzza, "Accurate angular velocity estimation with an event camera," *IEEE Robot. Autom. Lett.*, vol. 2, no. 2, pp. 632–639, 2017.
- [56] A. Z. Zhu, Y. Chen, and K. Daniilidis, "Realtime time synchronized event-based stereo," in *Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 438–452.
- [57] A. Z. Zhu, D. Thakur, T. Ozaslan, B. Pfrommer, V. Kumar, and K. Daniilidis, "The multivehicle stereo event camera dataset: An event camera dataset for 3D perception," *IEEE Robot. Autom. Lett.*, vol. 3, no. 3, pp. 2032–2039, Jul. 2018.
- [58] A. Mitrokhin, C. Ye, C. Fermuller, Y. Aloimonos, and T. Delbruck, "EV-IMO: Motion segmentation dataset and learning pipeline for event cameras," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2019, pp. 6105–6112.
- [59] S. Klenk, J. Chui, N. Demmel, and D. Cremers, "TUM-VIE: The TUM stereo visual-inertial event dataset," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2021, pp. 8601–8608.
- [60] H. Rebecq, D. Gehrig, and D. Scaramuzza, "ESIM: an open event camera simulator," in *Conf. on Robotics Learning (CoRL)*, ser. Proc. Machine Learning Research, vol. 87. PMLR, 2018, pp. 969–982.
- [61] E. Mueggler, H. Rebecq, G. Gallego, T. Delbruck, and D. Scaramuzza, "The event-camera dataset and simulator: Event-based data for pose estimation, visual odometry, and SLAM," *Int. J. Robot. Research*, vol. 36, no. 2, pp. 142–149, 2017.
- [62] M. Muglikar, M. Gehrig, D. Gehrig, and D. Scaramuzza, "How to calibrate your event camera," in *IEEE Conf. Comput. Vis. Pattern Recog. Workshops (CVPRW)*, 2021, pp. 1403–1409.
- [63] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the KITTI vision benchmark suite," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2012, pp. 3354–3361.
- [64] C. Ye, A. Mitrokhin, C. Parameshwara, C. Fermüller, J. A. Yorke, and Y. Aloimonos, "Unsupervised learning of dense optical flow, depth and egomotion with event-based sensors," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst.* (IROS), 2020, pp. 5831–5838.
- [65] M. Pizzoli, C. Forster, and D. Scaramuzza, "REMODE: Probabilistic, monocular dense reconstruction in real time," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2014, pp. 2609–2616.
- [66] H. Hirschmuller, "Stereo processing by semiglobal matching and mutual information," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 2, pp. 328–341, Feb. 2008.
- [67] C. Brandli, T. Mantel, M. Hutter, M. Höpflinger, R. Berner, R. Siegwart, and T. Delbruck, "Adaptive pulsed laser line extraction for terrain reconstruction using a dynamic vision sensor," *Front. Neurosci.*, vol. 7, p. 275, 2014.
- [68] S. Tulyakov, F. Fleuret, M. Kiefel, P. Gehler, and M. Hirsch, "Learning an event sequence embedding for dense event-based deep stereo," in *Int. Conf. Comput. Vis. (ICCV)*, 2019, pp. 1527–1537.
- [69] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, "SVO: Semidirect visual odometry for monocular and multicamera systems," *IEEE Trans. Robot.*, vol. 33, no. 2, pp. 249– 265, 2017.
- [70] V. Usenko, N. Demmel, D. Schubert, J. Stueckler, and D. Cremers, "Visual-inertial mapping with non-linear factor recovery," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 422–429, 2020.
- [71] L. Burner, A. Mitrokhin, C. Fermüller, and Y. Aloimonos, "EVIMO2: An event camera dataset for motion segmentation, optical flow, structure from motion, and visual inertial odometry in indoor scenes with monocular or stereo algorithms," arXiv eprints, May 2022.
- [72] B. Son, Y. Suh, S. Kim, H. Jung, J.-S. Kim, C. Shin, K. Park, K. Lee, J. Park, J. Woo, Y. Roh, H. Lee, Y. Wang, I. Ovsiannikov, and H. Ryu, "A 640x480 dynamic vision sensor with a 9μm pixel and 300Meps address-event representation," in *IEEE Intl. Solid-State Circuits Conf. (ISSCC)*, 2017, pp. 66–67.
- [73] Prophesee Evaluation Kits, https://www.prophesee.ai/ event-based-evk/, 2020.
- [74] G. Taverni, D. P. Moeys, C. Li, C. Cavaco, V. Motsnyi, D. S. S. Bello, and T. Delbruck, "Front and back illuminated Dynamic and Active Pixel Vision Sensors comparison," *IEEE Trans. Circuits Syst. II*, vol. 65, no. 5, pp. 677–681, 2018.
- [75] P. Dudek, T. Richardson, L. Bose, S. Carey, J. Chen, C. Greatwood, Y. Liu, and W. Mayol-Cuevas, "Sensor-level computer vision with pixel processor arrays for agile robots," *Science Robotics*, vol. 7, no. 67, p. eabl7755, 2022.
- [76] A. Z. Zhu, L. Yuan, K. Chaney, and K. Daniilidis, "EV-FlowNet: Self-supervised optical flow estimation for event-based cameras," in *Robotics: Science and Systems (RSS)*, 2018, pp. 1–9.

- [77] J. Hidalgo-Carrió, G. Gallego, and D. Scaramuzza, "Event-aided direct sparse odometry," in *IEEE Conf. Comput. Vis. Pattern Recog.* (CVPR), Jun. 2022, pp. 5781–5790.
- [78] Z. Wang, D. Yuan, Y. Ng, and R. Mahony, "A linear comb filter for event flicker removal," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2022, pp. 398–404.
- [79] A. J. Davison, "FutureMapping: The computational structure of spatial AI systems," arXiv e-prints, Mar. 2018.



Suman Ghosh is pursuing doctoral studies in the Department of Electrical Engineering and Computer Science at Technische Universität Berlin, Germany. Before that, he worked as Research Fellow at the EDPR group at Istituto Italiano di Tecnologia, Italy. He obtained his Masters degree on Robotics in 2019 from the University of Genoa, Italy, and Warsaw University of Technology, Poland, with full Erasmus Mundus scholarship under the European Master on Advanced Robotics (EMARO+) program. In 2016,

he received a Bachelors degree in Electronics and Telecommunication Engineering from Jadavpur University, India. His research interests include (biological and computer) vision, robotics and learning.



Guillermo Gallego is Associate Professor at Technische Universität Berlin, Berlin, Germany, in the Dept. of Electrical Engineering and Computer Science, and at the Einstein Center Digital Future, Berlin, Germany. He is also a Principal Investigator at the Science of Intelligence Excellence Cluster, Berlin, Germany. He received the PhD degree in Electrical and Computer Engineering from the Georgia Institute of Technology, USA, in 2011, supported by a Fulbright Scholarship. From 2011 to 2014 he was a Marie Curie

researcher with Universidad Politecnica de Madrid, Spain, and from 2014 to 2019 he was a postdoctoral researcher at the Robotics and Perception Group, University of Zurich and ETH Zurich, Switzerland. He serves as Associate Editor for IEEE Transactions on Pattern Analysis and Machine Intelligence, and for IEEE Robotics and Automation Letters. His research interests include robotics, computer vision, signal processing, optimization and geometry.