

# End-to-End Learning of Representations for Asynchronous Event-Based Data

Daniel Gehrig<sup>1</sup>, Antonio Loquercio<sup>1</sup>, Konstantinos G. Derpanis<sup>2</sup>, Davide Scaramuzza<sup>1</sup>

<sup>1</sup> Robotics and Perception Group

Depts. Informatics and Neuroinformatics

University of Zurich and ETH Zurich

<sup>2</sup> Ryerson University and Samsung AI Centre Toronto

## Abstract

Event cameras are vision sensors that record asynchronous streams of per-pixel brightness changes, referred to as “events”. They have appealing advantages over frame-based cameras for computer vision, including high temporal resolution, high dynamic range, and no motion blur. Due to the sparse, non-uniform spatiotemporal layout of the event signal, pattern recognition algorithms typically aggregate events into a grid-based representation and subsequently process it by a standard vision pipeline, e.g., Convolutional Neural Network (CNN). In this work, we introduce a general framework to convert event streams into grid-based representations through a sequence of differentiable operations. Our framework comes with two main advantages: (i) allows learning the input event representation together with the task dedicated network in an end-to-end manner, and (ii) lays out a taxonomy that unifies the majority of extant event representations in the literature and identifies novel ones. Empirically, we show that our approach to learning the event representation end-to-end yields an improvement of approximately 12% on optical flow estimation and object recognition over state-of-the-art methods.

## Multimedia Material

The project’s code is available on the following page: [https://github.com/uzh-rpg/rpg\\_event\\_representation\\_learning](https://github.com/uzh-rpg/rpg_event_representation_learning). Additionally, qualitative results can be viewed in this video: <https://youtu.be/bQtSx59GXRY>

## 1. Introduction

Event cameras are bio-inspired vision sensors that operate radically differently from traditional cameras. Instead of capturing brightness images at a fixed rate, event cameras measure brightness *changes* (called *events*) for each pixel independently. Event cameras, such as the Dy-

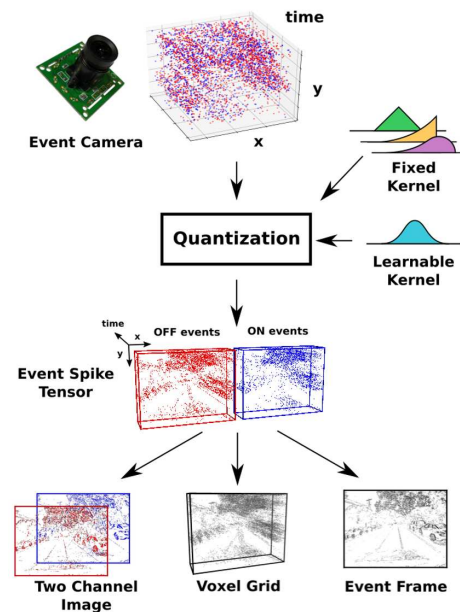


Figure 1. General framework to convert asynchronous event data into grid-based representations using convolutions, quantization, and projections. All of these operations are differentiable. Best viewed in color.

amic Vision Sensor (DVS) [34], possess appealing properties compared to traditional frame-based cameras, including a very high dynamic range, high temporal resolution (in the order of microseconds), and low power consumption. In addition, event cameras greatly reduce bandwidth. While frame-based cameras with comparable temporal resolution and/or dynamic range cameras exist, they are typically bulky, power-intensive, and require cooling [48].

The output of an event camera consists of a stream of events that encode the time, location, and polarity (sign) of the brightness changes. Consequently, each event alone carries very little information about the scene. Event-based vision algorithms aggregate information to enable further processing in two ways: (i) use a continuous-time model (e.g., Kalman filter) that can be updated asynchronously

with each incoming event [3, 19, 29, 40, 44] or (ii) process events simultaneously in packets [27, 52, 57, 66, 67], *i.e.*, spatiotemporal localized aggregates of events. The former methods can achieve minimal latency, but are sensitive to parameter tuning (*e.g.*, filter weights) and are computationally intensive, since they perform an update step for each event. In contrast, methods operating on event packets trade-off latency for computational efficiency and performance. Despite their differences, both paradigms have been successfully applied on various vision tasks, including tracking [19, 21, 40, 42], depth estimation [3, 52, 67], visual odometry [27, 54, 57, 66], recognition [29, 44], and optical flow estimation [7, 69]. A good survey on the applications of event cameras can be found in [18]. Motivated by the broad success of deep learning in computer vision on frame-based imagery, a growing number of recent event-based works have adopted a data driven approach [2, 32, 36, 47, 69]. Spiking Neural Networks (SNNs) are a natural fit to process event streams, since they enable asynchronous inference at low power on specialized hardware [2, 32, 47]. However, SNNs are notoriously difficult to train, as no efficient backpropagation algorithm exists [24]. In addition, the special-purpose hardware required to run SNNs is expensive and in the development stage, which hinders its widespread adoption in the vision community.

Most closely related to the current paper are methods that pair an event stream with standard frame-based deep convolutional neural network (CNN) or recursive architectures, *e.g.*, [29, 36, 41, 59, 69]. To do so, a pre-processing step typically converts asynchronous event data to a grid-like representation, which can be updated either synchronously [36, 69] or asynchronously [29, 59]. These methods benefit from their ease of implementation using standard frame-based deep learning libraries (*e.g.*, [1, 46]) and fast inference on commodity graphics hardware. However, these efforts have mainly focused on the downstream task beyond the initial representational stage and simply consider a fixed, possibly suboptimal, conversion between the raw event stream and the input grid-based tensor. To date, there has not been an extensive study on the impact of the choice of input representation, leaving the following fundamental open question: *What is the best way to convert an asynchronous event stream into a grid-based (tensor) representation to maximize the performance on a given task?* In this paper, we aim to address this knowledge gap.

**Contributions** We propose a general framework that converts asynchronous event-based data into grid-based representations. To achieve this, we express the conversion process through kernel convolutions, quantizations, and projections, where each operation is differentiable (see Fig. 1). Our framework comes with two main advantages. First, it makes the conversion process fully differentiable, allowing to learn a representation *end-to-end* from raw event data to

the task loss. In contrast, prior work assumes the input event representation as fixed. Second, it lays out a taxonomy that unifies the majority of extant event representations in the literature and identifies novel ones. Through extensive empirical evaluations we show that our approach to learning the event representation end-to-end yields an improvement of 12% on optical flow and 12.6% on object recognition over state-of-the-art approaches that rely on handcrafted input event representations. In addition, we compare our methodology to asynchronous approaches in term of accuracy and computational load to shed light on the relative merits of each category.

## 2. Related Work

Traditionally, handcrafted features were used in frame-based computer vision, *e.g.*, [14, 33, 35, 60, 63]. More recently, research has shifted towards data-driven models, where features are automatically learned from data, *e.g.*, [4, 22, 25, 37, 55]. The main catalyst behind this paradigm shift has been the availability of large training datasets [12, 15, 16], efficient learning algorithms [30, 61] and suitable hardware. Only recently has event-based vision made strides to address each of these areas.

Analogous to early frame-based computer vision approaches, significant effort has been made in designing efficient spatiotemporal feature descriptors of the event stream. From this line of research, typical high-level applications are gesture recognition [31], object recognition [29, 45, 59] or face detection [6]. Low-level applications include optical flow prediction [8, 9] and image reconstruction [5].

Another line of research has focused on applying data-driven models to event-based data. These include asynchronous, spiking neural networks (SNNs)<sup>1</sup> [32], which have been applied to several tasks, *e.g.*, object recognition [32, 44, 47, 64], gesture classification [2], and optical flow prediction [7, 8]. However, the lack of specialized hardware and computationally efficient backpropagation algorithms still limits the usability of SNNs in complex real-world scenarios. A typical solution to this problem is learning parameters with frame-based data and transferring the learned parameters to event data [13, 47]. However, it is not clear how much this solution can generalize to real, noisy, event data that has not been observed during training.

Recently, several works have proposed to use standard learning architectures as an alternative to SNNs [36, 41, 59, 68, 69]. To process asynchronous event streams, Neil et al. [41] adapted a recursive architecture to include the time dimension for prediction. Despite operating asynchronously, their approach introduces high latency, since

<sup>1</sup>Here we use the term SNNs as in the neuromorphic literature [32], where it describes continuous-time neural networks. Other networks which are sometimes called SNNs are low precision networks, such as binary networks [51]. However, these are not well suited for asynchronous inputs.

Representation	Dimensions	Description	Characteristics
Event frame [53]	$H \times W$	Image of event polarities	Discards temporal and polarity information
Event count image [36, 69]	$2 \times H \times W$	Image of event counts	Discards time stamps
Surface of Active Events (SAE) [7, 69]	$2 \times H \times W$	Image of most recent time stamp	Discards earlier time stamps
Voxel grid [70]	$B \times H \times W$	Voxel grid summing event polarities	Discards event polarity
Histogram of Time Surfaces (HATS) [59]	$2 \times H \times W$	Histogram of average time surfaces	Discards temporal information
Event Spike Tensor (EST, our work)	$2 \times B \times H \times W$	Sample event point-set into a grid	Discards the least amount of information

Table 1. Comparison of grid-based event representations used in prior work on event-based deep learning.  $H$  and  $W$  denote the image height and width dimensions, respectively, and  $B$  the number of temporal bins.

events have to pass sequentially through the entire recursive structure. To reduce latency, other methods convert event streams into a grid-based representation, compatible with learning algorithms designed for standard frames, *e.g.*, CNNs [36, 59, 68, 69]. Sironi et al. [59] obtained state-of-the-art results in object recognition tasks by transforming events into histograms of averaged time surfaces (HATS), which are then fed to a support vector machine for inference. The main advantage of their representation is that it can not only be used in conjunction with standard learning pipelines, but it can also be updated asynchronously, if sufficient compute is available. A simpler representation was proposed by Maqueda et al. [36] to address steering-angle prediction, where events of different polarities are accumulated over a constant temporal window. To perform a low-level task, *i.e.*, optical flow estimation, Zhu et al. [69] proposed to convert events into a four-dimensional grid that includes both the polarity and spike time. Finally, Zhu et al. [70] converted events into a spatiotemporal voxel-grid. Compared to the representation proposed in [36], the two latter representations have the advantage of preserving temporal information. A common aspect among these works is the use of a handcrafted event stream representation. In contrast, in this paper we propose a novel event-based representation that is learned end-to-end together with the task. A comparison of event-based representations and their design choices is summarized in Table 1.

Coupling event-based data with standard frame-based learning architectures has the potential to realize the flexibility of learning algorithms with the advantages of event cameras. It is however not yet clear what is the impact of the event representation on the task performance. In this work, we present an extensive empirical study on the choice of representation for the tasks of object recognition and optical flow estimation, central tasks in computer vision.

### 3. Method

In this section, we present a general framework to convert asynchronous event streams into grid-based representations. By performing the conversion strictly through differentiable operators, our framework allows us to learn a representation end-to-end for a given task. Equipped with this tool, we derive a taxonomy that unifies common representations in the literature and identifies new ones. An overview of the proposed framework is given in Fig. 2.

#### 3.1. Event Data

Event cameras have pixels which trigger events independently whenever there is a log brightness change:

$$L(x, y, t) - L(x, y, t - \Delta t) \geq pC, \quad (1)$$

where  $C$  is the contrast threshold,  $p \in \{-1, 1\}$  is the polarity of the change in brightness, and  $\Delta t$  is the time since the last event at  $\mathbf{u} = (x, y)^\top$ . In a given time interval  $\Delta\tau$ , the event camera will trigger a number of events:

$$\mathcal{E} = \{e_k\}_{k=1}^N = \{(x_k, y_k, t_k, p_k)\}_{k=1}^N. \quad (2)$$

Due to their asynchronous nature, events are represented as a set. To use events in combination with a convolutional neural network it is necessary to convert the event set into a grid-like representation. This means we must find a mapping  $\mathcal{M} : \mathcal{E} \mapsto \mathcal{T}$  between the set  $\mathcal{E}$  and a tensor  $\mathcal{T}$ . Ideally, this mapping should preserve the structure (*i.e.*, spatiotemporal locality) and information of the events.

#### 3.2. Event Field

Intuitively, events represent point-sets in a four-dimensional manifold spanned by the  $x$  and  $y$  spatial coordinates, time, and polarity. This point-set can be summarized by the *event field*, inspired by [11, 32]:

$$S_{\pm}(x, y, t) = \sum_{e_k \in \mathcal{E}_{\pm}} \delta(x - x_k, y - y_k) \delta(t - t_k), \quad (3)$$

defined in continuous space and time, for events of positive ( $\mathcal{E}_+$ ) and negative ( $\mathcal{E}_-$ ) polarity. This representation replaces each event by a Dirac pulse in the space-time manifold. The resulting function  $S_{\pm}(x, y, t)$  gives a continuous-time representation of  $\mathcal{E}$  which preserves the event’s high temporal resolution and enforces spatiotemporal locality.

#### 3.3. Generating Representations

**Measurements** In this section, we generalize the notion of the event field and demonstrate how it can be used to generate a grid-like representation from the events. We observe that (3) can be interpreted as successive measurements of a function  $f_{\pm}$  defined on the domain of the events, *i.e.*,

$$S_{\pm}(x, y, t) = \sum_{e_k \in \mathcal{E}_{\pm}} f_{\pm}(x, y, t) \delta(x - x_k, y - y_k) \delta(t - t_k). \quad (4)$$

We call (4) the *Event Measurement Field*. It assigns a measurement  $f_{\pm}(x_k, y_k, t_k)$  to each event. Examples of such functions are the event polarity  $f_{\pm}(x, y, t) = \pm 1$ , the event count  $f_{\pm}(x, y, t) = 1$ , and the normalized time stamp  $f_{\pm}(x, y, t) = \frac{t-t_0}{\Delta t}$ . Other examples might include the instantaneous event rate or image intensity provided by such sensors as the Asynchronous Time-based Image Sensor (ATIS) [10]. Various representations in the literature make use of the event measurement field. In several works [29, 36, 59, 69], pure event counts are measured, and summed for each pixel and polarity to generate *event count images*. Other works [7, 69] use the time stamps of the events to construct the *surface of active events* (SAE) which retains the time stamp of the most recent event for each pixel and polarity. Other representations use the event polarities and aggregate them into a three-dimensional *Voxel Grid* [70] or a two-dimensional *Event Frame* [53].

**Kernel Convolutions** Although the event measurement field retains the high temporal resolution of the events, it is still ill-defined due to the use of Dirac pulses. Therefore, to derive a meaningful signal from the event measurement field, we must convolve it with a suitable aggregation kernel. The convolved signal thus becomes:

$$(k * S_{\pm})(x, y, t) = \sum_{e_k \in \mathcal{E}_{\pm}} f_{\pm}(x_k, y_k, t_k) k(x - x_k, y - y_k, t - t_k). \quad (5)$$

In the literature, (5) is also known as the *membrane potential* [32, 39, 49]. Several variations of this kernel have been used in prior works. The two most commonly used ones are the *alpha-kernel*,  $k(x, y, t) = \delta(x, y) \frac{e^{-t/\tau}}{\tau} \exp(-t/\tau)$  [32, 39], and the exponential kernel,  $k(x, y, t) = \delta(x, y) \frac{1}{\tau} \exp(-t/\tau)$  [49]. In fact, the exponential kernel is also used to construct the hierarchy of time-surfaces (HOTS) [29] and histogram of average time-surfaces (HATS) [59], where events are aggregated into *exponential time surfaces*. In the case of HATS [59], the exponential time surfaces can be interpreted as a local convolution of the spike train with an exponential kernel. Another kernel which is typically used is the trilinear voting kernel,  $k(x, y, t) = \delta(x, y) \max(0, 1 - |\frac{t}{\Delta t}|)$  [26]. Generally, the design of kernel functions is based on task-dependent heuristics with no general agreement on the optimal kernel to maximize task performance.

**Discretized Event Spike Tensor** After kernel convolutions, a grid representation of events can be realized by sampling the convolved signal, (5), at regular intervals:

$$S_{\pm}[x_l, y_m, t_n] = (k * S_{\pm})(x_l, y_m, t_n) = \sum_{e_k \in \mathcal{E}_{\pm}} f_{\pm}(x_k, y_k, t_k) k(x_l - x_k, y_m - y_k, t_n - t_k). \quad (6)$$

Typically, the spatiotemporal coordinates,  $x_l, y_m, t_n$ , lie on a voxel grid, *i.e.*,  $x_l \in \{0, 1, \dots, W-1\}$ ,  $y_m \in \{0, 1, \dots, H-1\}$ , and  $t_n \in \{t_0, t_0 + \Delta t, \dots, t_0 + B\Delta t\}$ , where  $t_0$  is the first time stamp,  $\Delta t$  is the bin size, and  $B$  is the number of temporal bins. We term this generalized representation the *Event Spike Tensor* (EST). Summing over both the polarity and time dimensions, one can derive the event-frame representation introduced in prior work [57]. Previous works considered quantizing various dimensions, including spatiotemporal binning [70], and quantizing both the polarity and spatial dimensions [36, 69]. However, the generalized form that retains all four dimensions has not been previously considered, and thus is a *new representation*.

**End-to-end Learned Representations** The measurement and kernel in (6) are generally hand crafted functions. Previous works manually tuned those functions to maximize task performance. In contrast, we propose to leverage the data directly to find the best function candidate, thus learning the representation *end-to-end*. We achieve this by replacing the kernel function in (6) with a multilayer perceptron (MLP) with two hidden layers each with 30 units. This MLP takes the coordinates and time stamp of an event as input, and produces an activation map around it. For each grid location in the representation we evaluate the activation maps produced by each event and sum them together according to (6). This operation is repeated for every point in the final grid, resulting in a grid-like representation. To enforce symmetry across events, we limit the MLP input to the difference in coordinates  $x_l - x_k, y_m - y_k, t_l - t_k$ . For the sake of simplicity, we do not learn the measurement function as well, choosing it instead from a set of fixed functions. To speed up inference, at test time the learnt kernel can be substituted with an efficient look-up table, thus having comparable computation cost to handcrafted kernels. These design choices make the representation both efficient and fully differentiable. In contrast to previous works that used sub-optimal heuristics to convert events into grids, our framework can now tune the representation to the downstream task, thus maximizing performance.

**Projection** From the generalized *event spike tensor* we can further instantiate novel and existing representations. Many works for example deal with three-dimensional tensors, such as [29, 36, 59, 69, 70]. The event spike tensor, being a four-dimensional data structure (two spatial, one temporal, and one polarity), thus acts as a precursor to these three-dimensional constructs, which can be obtained by summing over one of the four dimensions. For example, the Two-Channel Image [36, 59, 69], can be derived by contracting the temporal dimension, either through summation [36, 59, 69] or maximization [69]. The voxel grid representation [70] can be derived by summing across event polarities. All of these operations can be generalized via



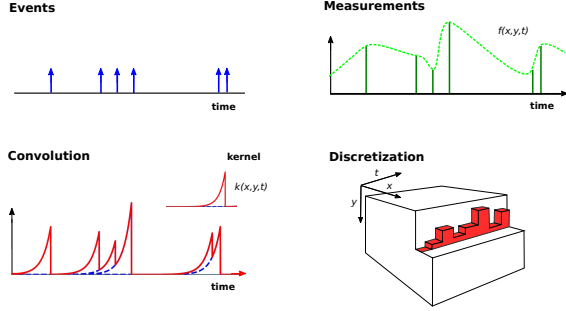


Figure 2. An overview of our proposed framework. Each event is associated with a measurement (green) which is convolved with a (possibly learnt) kernel. This convolved signal is then sampled on a regular grid. Finally, various representations can be instantiated by performing projections over the temporal axis or over polarities.

the projection operator  $H_v$ , where  $H$  can be summation  $\Sigma$ , maximization  $\max$ , *etc.* and  $v$  denoting the dimension can be  $x_l, y_m, t_n$ , or over polarity  $\pm$ , yielding 16 possible projections. Here, we list only the representations that retain the spatial dimension, of which there are four, including the EST without projection:

$$S_{\pm}[x_l, y_m, t_n] \quad (7)$$

$$S[x_l, y_m, t_n] = H_{\pm}(S_{\pm}[x_l, y_m, t_n]) \quad (8)$$

$$S_{\pm}[x_l, y_m] = H_{t_n}(S_{\pm}[x_l, y_m, t_n]) \quad (9)$$

$$S[x_l, y_m] = H_{t_n, \pm}(S_{\pm}[x_l, y_m, t_n]). \quad (10)$$

We refer to these representations as the *EST* (7), *Voxel Grid* (8), *Two-Channel Image* (9), and *Event Frame* (10). The direction of projection has an impact on the information content of the resulting representation. For example, projecting along the temporal axis greatly compresses the event representation, but at the cost of temporal localization information. In contrast, projecting the event polarities leads to the cancellation of positive and negative events, potentially removing information in the process. Of these representations, the EST stands out, as it retains the maximum amount of event information by forgoing the projection operation.

## 4. Empirical Evaluation

In this section, we present an extensive comparative evaluation of the representations identified by our taxonomy for object recognition (Sec. 4.1) and optical flow estimation (Sec. 4.2) on standard event camera benchmarks.

**Candidate Representations** We start out by identifying 12 distinct representations based on the event spike tensor (6). In particular, we select the measurement function (4) from three candidates: event polarity, event count, and normalized time stamp. We use the summation operator  $\Sigma$  to project out various axes defined in (7) - (10), resulting

in four variations: Event Spike Tensor, Voxel Grid, Two-Channel Image, and Event Frame. We split the event spike tensor (a four-dimensional tensor) along the polarity dimension and concatenate the two tensors along the temporal dimension, effectively doubling the number of channels. This is done to make the representation compatible with two-dimensional convolutions. As a first step we apply a generic trilinear kernel to convolve the event spike signal, and later study the effect of different kernels on performance when applied to the EST. Finally, we report results for our end-to-end trained variant that directly utilizes raw events.

### 4.1. Object Recognition

Object recognition with conventional cameras remains challenging due to their low dynamic range, high latency, and tendency to motion blur. In recent years, event-based classification has grown in popularity because it can address all these challenges.

In this section, we investigate the performance of the event representations proposed in Sec. 4 on the task of event-based object recognition. In particular, we aim to determine the relationship between the information content of the representation and classification accuracy. We show that our end-to-end learned representation significantly outperforms the state-of-the-art [59]. We use two publicly available datasets in our evaluation: N-Cars [59] (Neuromorphic-Cars) and N-Caltech101 [43]. N-Cars provides a benchmark for the binary task of car recognition in a scene. It contains 24,029 event samples of 100 ms length recorded by the ATIS event camera [50]. N-Caltech101 (Neuromorphic-Caltech101) is the event-based version of the popular Caltech101 dataset [17], and poses the task of multiclass recognition for event cameras. It contains 8,246 samples and 100 classes, which were recorded by placing an event camera on a motor and moving it in front of a screen projecting various samples from Caltech101.

**Implementation** We use a ResNet-34 architecture [22] for each dataset. The network is pretrained on color RGB images from ImageNet [58]. To account for the different number of input channels and output classes between the pre-trained model and ours, following approach [36]: we replace the first and last layer of the pre-trained model with random weights and then finetune all weights on the task. We train by optimizing the cross-entropy loss and use the ADAM optimizer [28] with an initial learning rate of  $1e-5$ , which we reduce by a factor of two every 10,000 iterations. We use a batch-size of 60 and 100 for N-Caltech101 and N-Cars, respectively.

**Results** The classification results are shown in Table 2. From the representations that we evaluated, the event spike tensor with time stamp measurements has the highest accuracy on the test set for both N-Cars and N-Caltech101.

Representation	Measurement	Kernel	N-Cars	N-Caltech101
Event Frame	polarity	trilinear	0.866	0.587
Two-Channel Image			0.830	0.711
Voxel Grid			0.865	0.785
<b>EST (Ours)</b>			0.868	0.789
Event Frame	count	trilinear	0.799	0.689
Two-Channel Image			0.861	0.713
Voxel Grid			0.827	0.756
<b>EST (Ours)</b>			0.863	0.784
Event Frame	time stamps	trilinear	0.890	0.690
Two-Channel Image			0.917	0.731
Voxel Grid			0.847	0.754
<b>EST (Ours)</b>			0.917	0.787
<b>EST (Ours)</b>	time stamps	alpha	0.911	0.739
		exponential	0.909	0.782
		learned	<b>0.925</b>	<b>0.817</b>

Table 2. Classification accuracy for all event representations using different measurement functions, as described in Sec. 4. For each representation the temporal dimension was discretized into nine bins. For the best performing representation (EST and time stamp measurements) we additionally report results for different kernel choices: trilinear [26], exponential [49], alpha kernels [32], as well as a learnable kernel.

Representation	Measurement	Kernel	N-Cars	N-Caltech101
H-First [44]	-	-	0.561	0.054
HOTS [29]			0.624	0.210
Gabor-SNN [59]			0.789	0.196
HATS [59]			0.902	0.642
HATS + ResNet-34			0.909	0.691
Two-Channel Image [36]	count	trilinear	0.861	0.713
Voxel Grid [70]	polarity		0.865	0.785
<b>EST (Ours)</b>	time stamps	trilinear	0.917	0.787
		learned	<b>0.925</b>	<b>0.817</b>

Table 3. Comparison of the classification accuracy for different baseline representations [36, 70] and state-of-the-art classification methods [29, 44, 59]. As an additional baseline we pair the best performing representation from previous work (HATS [59]) with a more powerful classification model (ResNet-34, used in this work) as the original numbers were reported using a linear SVM.

From these results we can make two conclusions. First, we observe that representations that separate polarity consistently outperform those that sum over polarities. Indeed, this trend is observed for all measurement functions: discarding the polarity information leads to a decrease in accuracy of up to 7%. Second, we see that representations that retain the temporal localization of events, *i.e.*, the Voxel Grid and EST, consistently outperform their counterparts, which sum over the temporal dimension. These observations indicate that both polarity and temporal information are important for object classification. This trend explains why the EST leads to the most accurate predictions: it retains the maximum amount of information with respect to the raw event data.

Interestingly, using event time stamps as measurements is more beneficial than other measurements, since the information about polarity and event count is already encoded in the event spike tensor. Indeed, using the time stamps explicitly in the tensor partially recovers the high temporal resolution, which was lost during the convolution and discretization steps of the event field. We thus established that the EST with time stamp measurements performs best for

object classification. However, the effect of the temporal kernel remains to be explored. For this purpose we experimented with the kernels described in Sec. 3.3, namely the exponential [49], alpha [32], and trilinear [26] kernels. In addition, we evaluate our end-to-end trainable representation and report the results in Table 2. We see that using different handcrafted kernels negatively impacts the test accuracies. In fact, applying these kernels to the event spikes decreases the effective temporal localization compared to the trilinear kernel by overlapping the event signals in the representation. This makes it difficult for a network to learn efficiently how to identify individual events. Finally, we see that if we learn a kernel end-to-end we gain a significant boost in performance. This is justified by the fact that the learnable layer finds an optimal way to draw the events on a grid, maximizing the discriminativeness of the representation.

**Comparison with State-of-the-Art** We next compare our results with state-of-the-art object classification methods that utilize handcrafted event representations, such as HATS [59], HOTS [29], as well as a baseline implementation of an SNN [59]. For the best performing representation (HATS) we additionally report the classification accuracies obtained with the same ResNet-34 used to evaluate the EST; the original work used a linear SVM. Two additional baselines are used for comparison: (i) the histogram of events [36] (here Two-Channel Image), with event count measurements, and (ii) the Voxel Grid [70] with polarity measurements.

The results for these methods are summarized in Table 3. Our method outperforms the state-of-the-art (HATS) and variant (HATS + ResNet-34), as well as the Voxel Grid and Two-Channel Image baselines by 2.3%, 1.6%, 6% and 6.5% on N-Cars and 17.5%, 12.6%, 3.2% and 10.4% on N-Caltech101, respectively. In particular, we see that our representation is more suited for object classification than existing handcrafted features, such as HATS and HOTS, even if we use more complex classification models with these features. This is likely due to HATS discarding temporal information, which, as we established, plays an important role in object classification. It is important to note, compared to the state-of-the-art, our method does not operate asynchronously, or at low power with current hardware (as for example SNNs); however, we show in Sec. 4.3 that our method can still operate at a very high framerate that is sufficient for many high-speed applications.

## 4.2. Optical Flow Estimation

Like object recognition, optical flow estimation using frame-based methods remains challenging in high-dynamic range scenarios, *e.g.*, at night, and during high speed movements. In particular, motion blur and over/under-saturation of the sensor often violate brightness constancy in the

image, a fundamental assumption underlying many approaches, which leads to estimation errors. Due to their lack of motion blur and high dynamic range, event cameras have the potential to provide higher accuracy estimates in these conditions. Early works on event-based optical flow estimation fit planes to the spatiotemporal manifold generated by events [7]. Other works have tackled this task by finding the optimal event alignments when projected onto a frame [20, 65]. Most recently, the relatively large-scale Multi Vehicle Stereo Event Camera Dataset (MVSEC) [68] made possible deep learning-based optical flow [69, 70]. It provides data from a stereo DAVIS rig combined with a LIDAR for ground-truth optical flow estimation [69]. The dataset features several driving sequences during the day and night, and indoor sequences recorded onboard a quadcopter. The methods in [69, 70] learn flow in a self-supervised manner and use standard U-Net architectures [56], outperforming existing frame-based methods in challenging night-time scenarios. In [69], a four-channel image representation is used as input to the network. This image is comprised of the two-channel event count image used in [36] and two-channel surface of active events (SAE) [7], divided according to event polarities. While the event counts and time surfaces combine the temporal and spatial information of the event stream, it still compresses the event signal by discarding all event time stamps except the most recent ones.

To date, it is unclear which event representation is optimal to learn optical flow. We investigate this question by comparing the representations listed in Sec. 4 against the state-of-the-art [69] for the task of optical flow regression, evaluated on the MVSEC dataset.

**Implementation** We train an optical flow regressor on the outdoor sequences *outdoor\_day1* and *outdoor\_day2*. These sequences are split into about 40,000 samples at fixed time intervals. Each sample consists of events aggregated between two DAVIS frames, which are captured at 30 Hz. We use EV-FlowNet [69] as the base network, with the channel dimension of the initial convolution layer set to the same number of channels of each input representation. The network is trained from scratch using a supervised loss derived from ground truth motion field estimates:

$$l(f, f_{\text{gt}}) = \sum_x \rho(f - f_{\text{gt}}), \quad (11)$$

where  $\rho$  denotes the robust Charbonnier loss [62],  $\rho(x) = (x^2 + \epsilon^2)^\alpha$ . For our experiments, we chose  $\epsilon = 1e - 3$  and  $\alpha = 0.5$ . This loss is minimized using the ADAM optimizer [28] with an initial learning rate of  $5e - 5$  and reducing it by a factor of two after 40,000 iterations and then again every 20,000 iterations with a batch size of eight.

**Results** As in [69], we measure the performance of our networks by comparing the average end-point error (AEE =  $\frac{1}{N} \sum_i |f - f_{\text{gt}}|_2$ ) on the *indoor\_flying* datasets, which are

visually distinct from the training set. The test error on these datasets thus reflects the generalizability of our network, and its overall performance. In addition, as events only provide sparse information in the frame we only report the error computed at pixels where at least one event was triggered, as done in [69]. Following the KITTI 2015 benchmark [38], we report the percentage of pixels which have an end-point-error larger than three pixels and 5% of the ground-truth flow, also done in [69]. In the previous classification experiments we observed that time stamp measurements are essential for a discriminative representation. We thus focus on results obtained from representations using the time stamp as the measurement function, as well as different kernels. Table 4 summarizes the results obtained from this experiment. An exhaustive evaluation of the various measurement functions, *i.e.*, polarities and counts, as well as qualitative results, is available in the supplemental material.

From Table 4 we see that Voxel Grid and EST have similar AEE and outlier ratios. This indicates that optical flow estimation is not as sensitive to event polarity as observed for classification. This is further supported by the small performance gap between the Two-Channel image and Event Frame. A more striking difference comes when we compare representations which retain the temporal dimension (middle rows), with those that sum over it. Indeed, the accuracies of the Two-Channel Image and the Event Frame drop approximately 10 – 20% when compared to the EST and Voxel-Grid. As with the classification evaluation, we establish that EST is among the most competitive representations and further explore the influence of different kernels on the performance. These are summarized in the bottom set of rows of Table 4. We see that the exponential and alpha kernels outperform the trilinear kernel. This indicates a strong dependency on the kernel shape and we thus proceed with the fully end-to-end learnable version. As with classification, we observe that the learnable kernel significantly improves the accuracy on almost all scenes. The most significant improvements are achieved for outlier ratios, indicating that using learnable kernels improves the robustness of the system.

**Comparison with State-of-the-Art** We compare our method with the state-of-the-art [69], as well as other baselines based on the representations used in [36] and [70]. Table 4 presents a detailed comparison. It is clear that the EST outperforms the state-of-the-art by a large margin (12%). There are also significant improvements in terms of outlier ratio, reducing the outliers by an average of 49% which again indicates the robustness of our method. This performance difference is likely due the data-driven nature of the learnable EST. While existing approaches learn the task on fixed event representations, our method learns the task and representation jointly. The resulting representation is more adapted to the task and thus maximizes performance.

Representation	Measurement	Kernel	indoor_flying1		indoor_flying2		indoor_flying3	
			AEE	% Outlier	AEE	% Outlier	AEE	% Outlier
Two-Channel Image [36]	count	trilinear	1.21	4.49	2.03	22.8	1.84	17.7
EV-FlowNet [69]	-		1.03	2.20	1.72	15.1	1.53	11.9
Voxel Grid [70]	polarity		0.96	1.47	1.65	14.6	1.45	11.4
Event Frame	time stamps	trilinear	1.17	2.44	1.93	18.9	1.74	15.5
Two-Channel Image			1.17	1.5	1.97	14.9	1.78	11.7
Voxel Grid			0.98	1.20	1.70	14.3	1.5	12.0
EST (Ours)	time stamps	trilinear	1.00	1.35	1.71	11.4	1.51	8.29
		alpha	1.03	1.34	1.52	11.7	1.41	8.32
		exponential	<b>0.96</b>	1.27	1.58	10.5	<b>1.40</b>	9.44
		learnt	0.97	<b>0.91</b>	<b>1.38</b>	<b>8.20</b>	1.43	<b>6.47</b>

Table 4. Average end-point error (AEE) and % of outliers evaluation on the MVSEC dataset for different variations of the EST with time stamp measurements. For each representation the temporal dimension was discretized into nine bins. Various baselines [36, 70] and state-of-the-art methods [69] are compared.

### 4.3. Computational Time and Latency

One of the key advantages of event cameras are their low latency and high update rate. To achieve high-frequency predictions, previous works developed lightweight and fast algorithms to process each incoming event asynchronously. In contrast, other approaches aggregate events into packets and then process them simultaneously. While this sacrifices latency, it also leads to overall better accuracy, due to an increase in the signal-to-noise ratio. Indeed, in several pattern recognition applications, *e.g.*, object recognition and optical flow prediction, asynchronous processing is not essential: we may actually sacrifice it for improved accuracy. We compare these two modes of operation in Table 5 where we show the number of events that can be processed per second, as well as the total time used to process a single sample of 100 ms from the N-Cars dataset. It can be seen that if we allow for batch computation, our method using a learnt kernel and lookup table can run at a very high speed that is comparable to other methods. For applications where asynchronous updates or low-power consumption have higher priority than accuracy, other methods, *e.g.*, SNNs, hold an advantage with respect to our approach.

We further report the computation time per inference for different architectures in Table 6. We report the timing in two stages: representation computation and inference. While representation computation is performed on a CPU (Intel i7 CPU, 64bits, 2.7GHz and 16 GB of RAM), inference is performed on a GPU (GeForce RTX 2080 Ti). Table 6 shows that the computation of the representation only contributes a small part to the overall computation time, while most of the time is spent during inference. Nonetheless, we see that a full forward pass only takes on the order of 6 ms, which translates to a maximum inference rate of 146 Hz. Although not on the order of the event rate, this value is high enough for most high-speed applications, such as mobile robotics or autonomous vehicle navigation. Moreover, we see that we can reduce the inference time significantly if we use smaller models, achieving 255 Hz for a ResNet-18. Shallower models could potentially be run at minimal loss in accuracy by leveraging distillation techniques [23].

Method	Asynchronous	Time [ms]	Speed [kEv/s]
Gabor SNN [59]	Yes	285.95	14.15
HOTS [29]	Yes	157.57	25.68
HATS [59]	Yes	7.28	555.74
<b>EST (Ours)</b>	No	<b>6.26</b>	<b>632.9</b>

Table 5. Computation time for 100 ms of event data and number of events processed per second.

Model	Inference [ms]	Representation [ms]	Total [ms]	Rate [Hz]
ResNet-18	3.87	0.38	4.25	235
ResNet-34	6.47	0.38	6.85	146
ResNet-50	9.14	0.38	9.52	105
EV-FlowNet	5.70	0.38	6.08	164

Table 6. Computation time split into EST generation (0.38 ms) and inference for several standard network architectures. Both ResNet-34 [22] and EV-FlowNet [69] allow processing at approximately 146 Hz which is sufficient for most high-speed applications.

## 5. Conclusions

This paper presented a general framework for converting asynchronous event data into grid-based representations. By representing the conversion process through differentiable operations, our framework allows learning input representations in a data-driven fashion. In addition, our framework lays out a taxonomy which unifies a large number of extant event representations and identifies new ones. Through an extensive evaluation we show that learning representations end-to-end together with the task yields an increase of about 12% in performance over state-of-the-art methods, for the tasks of object recognition and optical flow estimation. With this contribution, we combined the benefits of deep learning with event cameras, thus unlocking their outstanding properties to a wider community. As an interesting direction for future work, we plan to allow asynchronous updates by deploying recurrent architectures, similar to [41]: this will bridge the gap between synchronous and asynchronous approaches for event-based processing.

## Acknowledgements

This project was funded by the Swiss National Center of Competence Research (NCCR) Robotics, through the Swiss National Science Foundation, and the SNSF-ERC starting grant. K.G.D. is supported by a Canadian NSERC Discovery grant. K.G.D. contributed to this work in his personal capacity as an Associate Professor at Ryerson University.



## References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: A system for large-scale machine learning. In *USENIX Conference on Operating Systems Design and Implementation*, pages 265–283, 2016. 2
- [2] Arnon Amir, Brian Taba, David Berg, Timothy Melano, Jeffrey McKinstry, Carmelo Di Nolfo, Tapan Nayak, Alexander Andreopoulos, Guillaume Garreau, Marcela Mendoza, Jeff Kusnitz, Michael Debole, Steve Esser, Tobi Delbruck, Myron Flickner, and Dharmendra Modha. A low power, fully event-based gesture recognition system. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 7388–7397, July 2017. 2
- [3] Alexander Andreopoulos, Hirak J. Kashyap, Tapan K. Nayak, Arnon Amir, and Myron D. Flickner. A low power, high throughput, fully event-based stereo system. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 7532–7542, 2018. 2
- [4] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 5297–5307, June 2016. 2
- [5] Patrick Bardow, Andrew J. Davison, and Stefan Leutenegger. Simultaneous optical flow and intensity estimation from an event camera. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 884–892, 2016. 2
- [6] Souptik Barua, Yoshitaka Miyatani, and Ashok Veeraraghavan. Direct face detection and video reconstruction from event cameras. In *IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, pages 1–9, 2016. 2
- [7] Ryad Benosman, Charles Clercq, Xavier Lagorce, Sio-Hoi Ieng, and Chiara Bartolozzi. Event-based visual flow. *IEEE Trans. Neural Netw. Learn. Syst.*, 25(2):407–417, 2014. 2, 3, 4, 7
- [8] Ryad Benosman, Sio-Hoi Ieng, Charles Clercq, Chiara Bartolozzi, and Mandyam Srinivasan. Asynchronous frameless event-based optical flow. *Neural Netw.*, 27:32–37, 2012. 2
- [9] Ryad Benosman, Sio-Hoi Ieng, Paul Rogister, and Christoph Posch. Asynchronous event-based Hebbian epipolar geometry. *IEEE Trans. Neural Netw.*, 22(11):1723–1734, 2011. 2
- [10] Christian Brandli, Raphael Berner, Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. A 240x180 130dB 3us latency global shutter spatiotemporal vision sensor. *IEEE J. Solid-State Circuits*, 49(10):2333–2341, 2014. 4
- [11] Andrea Censi. Efficient neuromorphic optomotor heading regulation. In *2015 American Control Conference (ACC)*, pages 3854–3861, July 2015. 3
- [12] Jia Deng, Wei Dong, R. Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*. IEEE, June 2009. 2
- [13] Peter U. Diehl, Daniel Neil, Jonathan Binas, Matthew Cook, Shih-Chii Liu, and Michael Pfeiffer. Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing. In *Int. Joint Conf. Neural Netw. (IJCNN)*, volume 4, pages 2933–2940, July 2015. 2
- [14] Piotr Dollar, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian detection: An evaluation of the state of the art. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(4):743–761, Apr. 2012. 2
- [15] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazırbaş, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Int. Conf. Comput. Vis. (ICCV)*, pages 2758–2766, 2015. 2
- [16] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, Sept. 2009. 2
- [17] Li Fei-Fei, R. Fergus, and P. Perona. One-shot learning of object categories. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(4):594–611, April 2006. 5
- [18] Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *arXiv e-prints*, abs/1904.08405, 2019. 2
- [19] Guillermo Gallego, Jon E. A. Lund, Elias Mueggler, Henri Rebecq, Tobi Delbruck, and Davide Scaramuzza. Event-based, 6-DOF camera tracking from photometric depth maps. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(10):2402–2412, Oct. 2018. 2
- [20] Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 3867–3876, 2018. 7
- [21] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. Asynchronous, photometric feature tracking using events and frames. In *Eur. Conf. Comput. Vis. (ECCV)*, 2018. 2
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 770–778, June 2016. 2, 5, 8
- [23] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *NIPS Deep Learning Workshop*, 2014. 8
- [24] Dongsung Huh and Terrence J Sejnowski. Gradient descent for spiking neural networks. In *Conf. Neural Inf. Process. Syst. (NIPS)*, pages 1440–1450, 2018. 2
- [25] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, July 2017. 2
- [26] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In *Conf.*

- Neural Inf. Process. Syst. (NIPS)*, pages 2017–2025, 2015. 4, 6
- [27] Hanme Kim, Stefan Leutenegger, and Andrew J. Davison. Real-time 3D reconstruction and 6-DoF tracking with an event camera. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 349–364, 2016. 2
- [28] Diederik P. Kingma and Jimmy L. Ba. Adam: A method for stochastic optimization. *Int. Conf. on Learning Representations (ICLR)*, 2015. 5, 7
- [29] Xavier Lagorce, Garrick Orchard, Francesco Gallupi, Bertram E. Shi, and Ryad Benosman. HOTS: A hierarchy of event-based time-surfaces for pattern recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(7):1346–1359, July 2017. 2, 4, 6, 8
- [30] Quoc V Le, Jiquan Ngiam, Adam Coates, Abhik Lahiri, Bobby Prochnow, and Andrew Y Ng. On optimization methods for deep learning. In *International Conference on Machine Learning (ICML)*, pages 265–272. Omnipress, 2011. 2
- [31] Junhaeng Lee, Tobi Delbruck, Paul K. J. Park, Michael Pfeiffer, Chang-Woo Shin, Hyunsurk Ryu, and Byung Chang Kang. Live demonstration: Gesture-based remote control using stereo pair of dynamic vision sensors. In *IEEE Int. Symp. Circuits Syst. (ISCAS)*, 2012. 2
- [32] Jun Haeng Lee, Tobi Delbruck, and Michael Pfeiffer. Training deep spiking neural networks using backpropagation. *Front. Neurosci.*, 10:508, 2016. 2, 3, 4, 6
- [33] Stefan Leutenegger, Margarita Chli, and Roland Y. Siegwart. BRISK: Binary robust invariant scalable keypoints. In *Int. Conf. Comput. Vis. (ICCV)*, pages 2548–2555, Nov. 2011. 2
- [34] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. A  $128 \times 128$  120 dB  $15 \mu\text{s}$  latency asynchronous temporal contrast vision sensor. *IEEE J. Solid-State Circuits*, 43(2):566–576, 2008. 1
- [35] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.*, 60(2):91–110, Nov. 2004. 2
- [36] Ana I. Maqueda, Antonio Loquercio, Guillermo Gallego, Narciso García, and Davide Scaramuzza. Event-based vision meets deep learning on steering prediction for self-driving cars. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 5419–5427, 2018. 2, 3, 4, 5, 6, 7, 8
- [37] Simon Meister, Junhwa Hur, and Stefan Roth. UnFlow: Un-supervised learning of optical flow with a bidirectional census loss. In *AAAI Conf. Artificial Intell.*, 2018. 2
- [38] Moritz Menze, Christian Heipke, and Andreas Geiger. Joint 3D estimation of vehicles and scene flow. In *ISPRS Workshop on Image Sequence Analysis (ISA)*, 2015. 7
- [39] Ammar Mohammed, Stefan Schliebs, and Nikola Kasabov. SPAN: A neuron for precise-time spike pattern association. In Bao-Liang Lu, Liqing Zhang, and James Kwok, editors, *Neural Information Processing*, pages 718–725. Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. 4
- [40] Elias Mueggler, Basil Huber, and Davide Scaramuzza. Event-based, 6-DOF pose tracking for high-speed maneuvers. In *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, pages 2761–2768, 2014. 2
- [41] Daniel Neil, Michael Pfeiffer, and Shih-Chii Liu. Phased LSTM: Accelerating recurrent network training for long or event-based sequences. In *Conf. Neural Inf. Process. Syst. (NIPS)*, pages 3882–3890, 2016. 2, 8
- [42] Zhenjiang Ni, Cécile Pacoret, Ryad Benosman, Sio-Hoi Ieng, and Stéphane Régner. Asynchronous event-based high speed vision for microparticle tracking. *J. Microscopy*, 245(3):236–244, 2012. 2
- [43] Garrick Orchard, Ajinkya Jayawant, Gregory K. Cohen, and Nitish Thakor. Converting static image datasets to spiking neuromorphic datasets using saccades. *Front. Neurosci.*, 9:437, 2015. 5
- [44] Garrick Orchard, Cedric Meyer, Ralph Etienne-Cummings, Christoph Posch, Nitish Thakor, and Ryad Benosman. HFirst: A temporal approach to object recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(10):2028–2040, 2015. 2, 6
- [45] Paul K. J. Park, Kyoobin Lee, Jun Haeng Lee, Byungkon Kang, Chang-Woo Shin, Jooyeon Woo, Jun-Seok Kim, Yunjae Suh, Sungho Kim, Saber Moradi, Ogan Gurel, and Hyunsurk Ryu. Computationally efficient, real-time motion recognition based on bio-inspired visual and cognitive processing. In *IEEE International Conference on Image Processing (ICIP)*. IEEE, Sept. 2015. 2
- [46] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS-W*, 2017. 2
- [47] José A. Perez-Carrasco, Bo Zhao, Carmen Serrano, Begoña Acha, Teresa Serrano-Gotarredona, Shouchun Chen, and Bernabé Linares-Barranco. Mapping from frame-driven to frame-free event-driven vision systems by low-rate rate coding and coincidence processing—application to feedforward ConvNets. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(11):2706–2719, Nov. 2013. 2
- [48] <https://www.phantomhighspeed.com/>, 2019. 1
- [49] Filip Ponulak. ReSuMe – new supervised learning method for spiking neural networks. 2005. 4, 6
- [50] Christoph Posch, Daniel Matolin, and Rainer Wohlgenannt. A QVGA 143 dB dynamic range frame-free PWM image sensor with lossless pixel-level video compression and time-domain CDS. *IEEE J. Solid-State Circuits*, 46(1):259–275, Jan. 2011. 5
- [51] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. XNOR-Net: Imagenet classification using binary convolutional neural networks. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 525–542, 2016. 2
- [52] Henri Rebecq, Guillermo Gallego, Elias Mueggler, and Davide Scaramuzza. EMVS: Event-based multi-view stereo—3D reconstruction with an event camera in real-time. *Int. J. Comput. Vis.*, pages 1–21, Nov. 2017. 2
- [53] Henri Rebecq, Timo Horstschaefer, and Davide Scaramuzza. Real-time visual-inertial odometry for event cameras using keyframe-based nonlinear optimization. In *British Machine Vis. Conf. (BMVC)*, Sept. 2017. 3, 4
- [54] Henri Rebecq, Timo Horstschaefer, Guillermo Gallego, and Davide Scaramuzza. EVO: A geometric approach to event-

- based 6-DOF parallel tracking and mapping in real-time. *IEEE Robot. Autom. Lett.*, 2:593–600, 2017. [2](#)
- [55] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. [2](#)
- [56] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, 2015. [7](#)
- [57] Antoni Rosinol Vidal, Henri Rebecq, Timo Horstschaefer, and Davide Scaramuzza. Ultimate SLAM? Combining events, images, and IMU for robust visual SLAM in HDR and high speed scenarios. *IEEE Robot. Autom. Lett.*, 3(2):994–1001, Apr. 2018. [2, 4](#)
- [58] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Fei-Fei Li. ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, Apr. 2015. [5](#)
- [59] Amos Sironi, Manuele Brambilla, Nicolas Bourdis, Xavier Lagorce, and Ryad Benosman. HATS: Histograms of averaged time surfaces for robust event-based object classification. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1731–1740, 2018. [2, 3, 4, 5, 6, 8](#)
- [60] Josef Sivic and Andrew Zisserman. Efficient visual search of videos cast as text retrieval. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31(4):591–606, Apr. 2009. [2](#)
- [61] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014. [2](#)
- [62] Deqing Sun, Stefan Roth, and Michael J Black. A quantitative analysis of current practices in optical flow estimation and the principles behind them. *International Journal of Computer Vision*, 106(2):115–137, 2014. [7](#)
- [63] Paul Viola and Michael J. Jones. Robust real-time face detection. *Int. J. Comput. Vis.*, 57(2):137–154, May 2004. [2](#)
- [64] Bo Zhao, Ruoxi Ding, Shoushun Chen, Bernabe Linares-Barranco, and Huajin Tang. Feedforward categorization on AER motion events using cortex-like features in a spiking neural network. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(9):1963–1978, Sept. 2015. [2](#)
- [65] Alex Zihao Zhu, Nikolay Atanasov, and Kostas Daniilidis. Event-based feature tracking with probabilistic data association. In *IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 4465–4470, 2017. [7](#)
- [66] Alex Zihao Zhu, Nikolay Atanasov, and Kostas Daniilidis. Event-based visual inertial odometry. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 5816–5824, 2017. [2](#)
- [67] Alex Zihao Zhu, Yibo Chen, and Kostas Daniilidis. Realtime time synchronized event-based stereo. In *Eur. Conf. Comput. Vis. (ECCV)*, 2018. [2](#)
- [68] Alex Zihao Zhu, Dinesh Thakur, Tolga Ozaslan, Bernd Pfrommer, Vijay Kumar, and Kostas Daniilidis. The multivehicle stereo event camera dataset: An event camera dataset for 3D perception. *IEEE Robot. Autom. Lett.*, 3(3):2032–2039, July 2018. [2, 3, 7](#)
- [69] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. EV-FlowNet: Self-supervised optical flow estimation for event-based cameras. In *Robotics: Science and Systems (RSS)*, 2018. [2, 3, 4, 7, 8](#)
- [70] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Unsupervised event-based learning of optical flow, depth, and egomotion. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, June 2019. [3, 4, 6, 7, 8](#)