

Sensor Generalization for Adaptive Sensing in Event-based Object Detection via Joint Distribution Training

Aheli Saha, René Schuster, and Didier Stricker

Deutsches Forschungszentrum für Künstliche Intelligenz, Kaiserslautern, Germany
firstName.lastName@dfki.de

Keywords: Object Detection, Event Cameras, Adaptive Sensors, Domain Generalization

Abstract: Bio-inspired event cameras have recently attracted significant research due to their asynchronous and low-latency capabilities. These features provide a high dynamic range and significantly reduce motion blur. However, because of the novelty in the nature of their output signals, there is a gap in the variability of available data and a lack of extensive analysis of the parameters characterizing their signals. This paper addresses these issues by providing readers with an in-depth understanding of how intrinsic parameters affect the performance of a model trained on event data, specifically for object detection. We also use our findings to expand the capabilities of the downstream model towards sensor-agnostic robustness.

1 INTRODUCTION

Object detection is essential to scene comprehension and forms the basis of increasingly complex visual perception tasks. Though traditional frame-based object detection has been widely investigated over the years, several challenges remain in highly dynamic environments. Frame-based cameras capture the entire frame at fixed, regular time intervals. This discrete-time sampling leads to imprecise detection, as fast-moving or transient objects may be completely missed or appear blurry across frames, making it difficult to detect them. Frame-based cameras also experience substantial latency. Low perceptual latency and accurate and precise decision-making are, however, critical in applications such as autonomous driving. Additionally, static and unchanging components of the scene are captured in each frame, leading to redundant data and requiring high power and bandwidth consumption.

Event cameras have recently gained considerable significance in this regard. They respond to intensity changes at individual pixel positions, triggering an event each time the difference in log intensity crosses a threshold. The recorded events encode the pixel coordinates, timestamp of occurrence (in millisecond precision), and the binary polarity of the brightness change. This operational technique allows event cameras to achieve a high dynamic range (> 120 dB) and striking resistance to motion blur. Their asynchronous operation and temporal precision in milliseconds pro-

vide unmatched benefits for real-time perception in constantly evolving scenarios. Due to these developments, event-based object detection has become an expanding field of research (see [Section 2.1](#)), particularly in the area of autonomous systems, such as drones, vehicles, or robots. However, to maximize efficiency in size- or power-sensitive applications, it can be argued that sensing and perception need to be considered jointly ([Eckmann et al., 2020](#); [Sprague et al., 2015](#)), as is observed in biological organisms to counteract resource constraints. This brings us to the theory of active efficient coding observed in biological organisms, in which the dynamic adaptation of input signals, in conjunction with the environment and task requirements, plays a crucial role. To simulate this evolutionary energy-efficient behavior in artificial systems, we aim to design an adaptive sensor. The intrinsic and extrinsic characteristics of that sensor can be manipulated by feedback loops to provide optimal information for a certain task in a specific environment, in contrast to a fixed sensor transduction. Ultimately, this feedback loop is based on performance and energy statistics, which allows for dynamic adjustment for varying environmental stimuli or the current task requirements. A visual representation of this control loop is shown in [Fig. 1](#). Further detailed discussion of this feedback-based dynamic sensor control is beyond the scope of this paper.

However, a key hindrance in the training and functioning of this control loop is the dependence of the detector performance on static input characteristics.

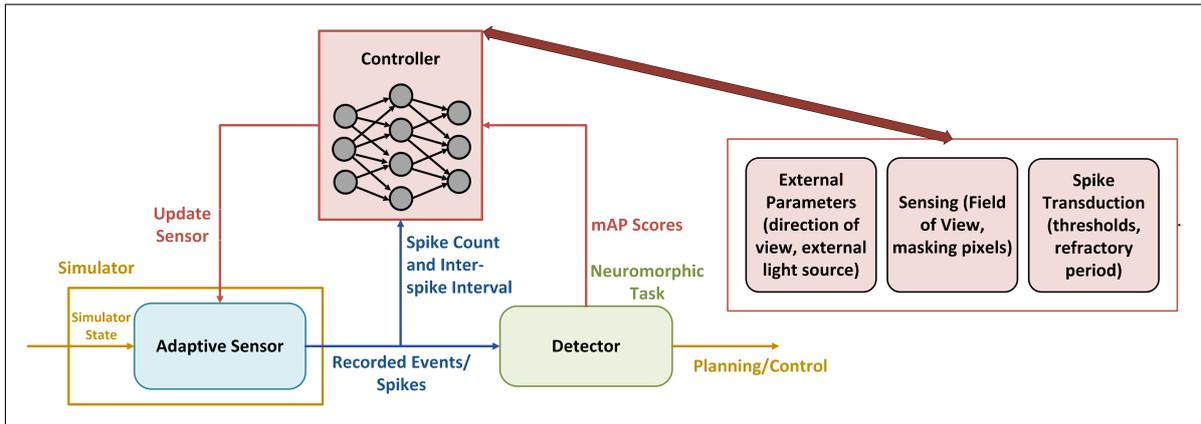


Figure 1: Overview of a sensor control loop. We aim to dynamically adapt the sensor characteristics based on energy utilization and task performance feedback. In this work, we investigate and mitigate the undesired effects of dynamically varying spike transduction on task performance, *i.e.*, object detection.

A sensor-agnostic detector is therefore a key requirement. In this work, we explore and investigate the challenge of changing event characteristics within an actively controlled sensory-perception system. Modifying the sensor characteristics alters the data distribution of the input signal, which a static, trained event-based object detector is unlikely to be capable of handling. To overcome this challenge, we employ a joint-training approach in which our object detector is trained on a set of data from a limited number of different sensors, effectively expanding the known input domain of our model. The contributions of our paper can be summarized as follows:

- We collect an expansive dataset of event data in simulation, with varying sensor characteristics (Section 3.2). To the best of our knowledge, this is the first dataset to probe this dimensionality of event data.
- We employ a well-known domain generalization training strategy to expand the model’s capabilities to interpolate across any possible sensor configuration in the parameter space.
- We design a rigorous exploratory framework and study and report the behavior of the model under distinct experimental test conditions, providing greater insights into the advances and limitations of this method.

2 RELATED WORK

2.1 Event-based Object Detection

Event-based object detection can be broadly classified into two categories (Silva et al., 2025) - (i) direct pro-

cessing of the sparse data, as in graph neural networks (GNNs) (Bi et al., 2019), and spiking neural networks (SNNs) (Cordone et al., 2022), and (ii) generation of dense event representations before application of traditional feature extractors.

GNNs use graphs with edges that encode the spatiotemporal connections between nearby events, represented by the nodes (Schaefer et al., 2022). This preserves the asynchronous, sparse character of event data while enabling localized updates; updates to the pertinent network subgraphs handle new events. While this significantly reduces computations, it is limited by the size and dimensionality of the local subgraphs (Gehrig and Scaramuzza, 2022).

The biologically inspired computation of SNNs is particularly well-suited for event data, due to their inherent compatibility with the asynchronous, sparse nature of the data. Each spiking neuron maintains an internal state and fires when a certain threshold is exceeded, at which stage the signal is propagated to connected neurons (Su et al., 2023; Yao et al., 2021; Zhang et al., 2022). SNNs maintain low latency and energy efficiency by avoiding static-frame processing and concentrating on dynamic changes. Compared to regular CNNs, SNNs use significantly less power since they only analyze active events- Spiking-YOLO requires approximately 280x lower energy than Tiny YOLO, (Kim et al., 2020; Su et al., 2023). However, SNNs require surrogate gradients (Bendig et al., 2023; Su et al., 2023; Guo et al., 2023; Cordone et al., 2022; Fan et al., 2024b) or ANN-to-SNN conversion (Bu et al., 2022; Deng and Gu, 2021; Wang et al., 2023; Li and Zeng, 2022), because of the non-differentiable nature of spike generation. Both of these methods lead to an efficiency-performance trade-off. Extensive research in this domain is still be-

ing conducted, with novel proposals addressing some existing challenges. EMS-YOLO (Su et al., 2023) enhanced the capacity to directly train deep-SNNs for object identification by creating a full-spike residual block. SFOD (Fan et al., 2024b) was the first to optimize SNNs for multi-scale feature integration.

Other branches aim to utilize dense feed-forward networks by composing a dense event representation from the sparse data. Some of the most common strategies are Event Histograms (Rebecq et al., 2017; Maqueda et al., 2018), Time surfaces (Sironi et al., 2018), Mixed Density Event Stacks (Nam et al., 2022), Time Ordered Event Representations (TORE) (Baldwin et al., 2022), and Event Temporal Images (Fan et al., 2024a). Here, we employ the Stacked Histogram Representations as in (Gehrig and Scaramuzza, 2023; Zubic et al., 2024) in which events in a temporal window are subdivided into smaller temporal bins, spatially distributed according to their pixel positions, and segregated into separate channels by polarity.

Significant progress was achieved with the introduction of benchmark large-scale datasets 1Mpx (Perot et al., 2020) and Gen1 (De Tournemire et al., 2020). They allowed for recent networks like RED (Perot et al., 2020) and ASTMNet (Li et al., 2022), which use memory methods to leverage the spatiotemporal information in event data fully. RVT (Gehrig and Scaramuzza, 2023), HMNet (Hamaguchi et al., 2023), and GET (Peng et al., 2023) have recently demonstrated outstanding performance employing Transformer networks on dense event representations.

2.2 Domain Generalization

A domain can be defined as the joint distribution between the input space and the labeled output space. Domain generalization (DG) refers to generalizing a model’s capabilities to out-of-distribution data that handles unseen test domains (Wang et al., 2022; Zhou et al., 2022). Of the two common and possible settings, single-source DG and multi-source DG, we are particularly interested in the latter, where we have access to multiple similar but distinct source domains. We aim to use this multi-source data to learn feature representations invariant to sensor settings of the source domains.

Current domain generalization research for event data focuses on the adaptability of models trained on limited synthetic data to real-world data, which, owing to non-alignment, results in performance degradation. The most prevalent source of disparity is the threshold determining the minimum per-pixel inten-

sity fluctuation necessary to trigger an event, which is susceptible to real-time shifts in a real camera (Planamente et al., 2021). Simulators fail to compensate for these non-idealities observed with actual cameras. To address this, some works like the ones by Stofregen et al. (2020) and Gehrig et al. (2020) operate at the input level, by simulating parameter variations. DA4Event (Planamente et al., 2021) proposes to bridge the sim-real gap at the feature level through a multi-view approach employing late fusion over multiple input event representations. Additionally, models trained for a particular frequency perform poorly at diverging inference frequencies (Zubic et al., 2024). To tackle this issue, FlexEvent (Lu et al., 2024) and FAOD (Zhang et al., 2024) propose modality-fusion and alignment techniques that can leverage both the semantic richness of RGB data and the fine-grained temporal resolution of event data. SSMS (Zubic et al., 2024) employs state-space models to train frequency-agnostic models.

Investigating another direction in generalized event cameras, Sundar et al. (2024) simulates a single-photon-based sensor design to demonstrate a conceptual design for event cameras capable of adjusting to spatio-temporal contexts. Although not investigated in this work, it is closely related to our final aim of producing an actively controlled event sensor, and our expanded model would play a significant role in the adaptability of such dynamic models.

Another closely related research direction is domain adaptation for event data, which addresses the scarcity of large-scale labeled event annotations, which impedes their inclusion into mainstream deep learning systems. CTN (Zhao et al., 2022) leverages the potential of both convolutional modules for early visual processing and Vision Transformers (ViTs) to capture long-range features. In DAEC² (Jian and Rostami, 2023), various augmentations of an object are projected to a latent space, and an encoder is trained to retain identity across these augmentations. However, in this work, we stick to domain generalization across multiple sensor configurations as our main goal, which is a mostly unexplored area in domain generalization, so far.

3 SENSOR-INVARIANT TRAINING

3.1 Data Collection

We use the CARLA simulator (Dosovitskiy et al., 2017) to collect our synthetic datasets. CARLA is pre-equipped with a Dynamic Vision Sensor (DVS)

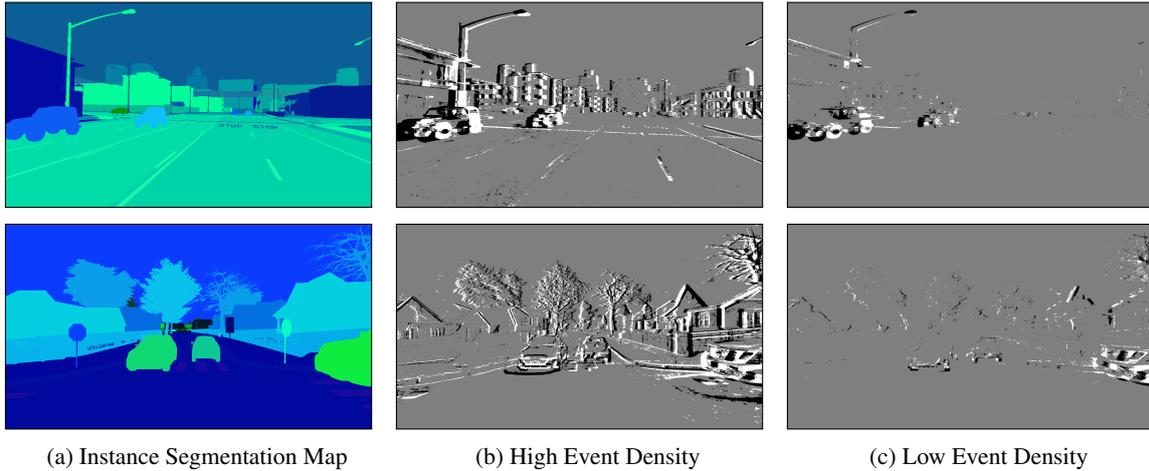


Figure 2: Examples of two scenes with varying configurations which produce high and low event densities, respectively. The black and white points represent negative and positive events, respectively, accumulated over a 50-ms time period. The gray spaces indicate the occurrence of no events within this timeframe at the corresponding pixel locations. The corresponding instance segmentation map is provided for reference to the actual vehicular locations.

camera, featuring an Application Programming Interface (API) to access its sensory characteristics, which is crucial in our case. Alongside the DVS sensor, we employ the RGB, depth, and instance segmentation sensors on our ego-vehicles to obtain accurate ground-truth bounding boxes for the objects visible in the scene. We retrieve the ground-truth boxes directly from the simulator, while the depth and segmentation modalities are employed to handle occluded or overlapping objects effectively. For demonstration, three classes—cars, buses, and trucks—have been merged into a single vehicle class to increase inter-class variability. Our employed methodology can be seamlessly expanded into multi-class scenarios.

However, the labels are obtained at a frequency corresponding to that of the RGB camera (20Hz), which is significantly lower than that of our event stream, which operates at a sub-millisecond latency. This is a limitation of the simulator. To overcome this, we use temporal windows to accumulate events into corresponding dense representations, such that each window coincides with the next set of labels from the simulator. Each temporal frame is further split into n_{bins} bins and segregated by polarity to produce the corresponding Stacked Histogram Representation, as by Gehrig and Scaramuzza (2023) and Zubic et al. (2024), with $(2 \cdot n_{bins} \times H \times W)$ tensor dimensions. Each of these tensors corresponds to one labeled output frame.

We collect data by driving our ego-vehicle across 12 distinct pre-determined routes in each of 13 different town maps, to produce variability in traffic density and weather conditions. 500 frames are collected for each sequence, at a frame rate of 20Hz, providing

us with 65 minutes of recorded data. Note that this frame rate refers to that of the traditional sensors, and only our event frames are synchronized to these. All events triggered in the corresponding 0.05s temporal window are recorded with their accurate and precise timestamps in *nanoseconds*. This produces high-quality, high-resolution (720 x 1280) data on driving environments, spread across diverse environmental scenarios and sensor characteristics. We also filter out bounding boxes with a side length lower than 20 pixels or a diagonal length below 60 pixels, following the evaluation criteria of Perot et al. (2020). Contrary to the 1Mpx dataset, we do not downsample our data and use the original high resolution. Towns are split for train, validation, and test, using a 70-15-15% split, resulting in (9, 2, 2) towns for each, respectively.

This collection procedure is repeated over 14 distinct DVS-sensor settings, outlined in Table 1 and discussed further in Section 3.2. Our dataset, of ~ 15 hours in total, thus provides a significant benchmark for numerous baseline and comparative studies on sensor behavior, enabling robust and sensor-invariant models for downstream tasks, and the study and development of dynamically adaptive event sensors.

3.2 Experimental Setup

3.2.1 Parameter Space

We focus on four key parameters that significantly influence event camera behavior - the intensity thresholds for event triggering (one for each polarity), the minimum threshold for interspike interval, and the

Table 1: Data Configurations.

| Config. (\mathcal{E}_i) | Pos. Thr. (th_p) | Neg. Thr. (th_n) | Ref. Period (T_r) [ms] | Field of View (F_v) |
|-----------------------------|----------------------|----------------------|----------------------------|-------------------------|
| \mathcal{E}_{base} | 0.5 | 0.5 | 10 | 90° |
| \mathcal{E}_1 | 0.25 | 0.25 | 0.01 | 90° |
| \mathcal{E}_2 | 0.75 | 0.75 | 0.01 | 90° |
| \mathcal{E}_3 | 1.0 | 1.0 | 0.01 | 90° |
| \mathcal{E}_4 | 0.5 | 0.5 | 10 | 90° |
| \mathcal{E}_5 | 0.5 | 0.5 | 25 | 90° |
| \mathcal{E}_6 | 0.5 | 0.5 | 50 | 90° |
| \mathcal{E}_7 | 0.5 | 0.5 | 0.01 | 45° |
| \mathcal{E}_8 | 0.5 | 0.5 | 0.01 | 135° |
| \mathcal{E}_9 | 0.5 | 0.5 | 0.01 | 160° |
| \mathcal{E}_{10} | 0.25 | 0.25 | 50 | 45° |
| \mathcal{E}_{11} | 1.0 | 0.5 | 25 | 90° |
| \mathcal{E}_{12} | 0.7 | 0.7 | 20 | 65° |
| \mathcal{E}_{13} | 0.3 | 0.9 | 15 | 130° |

field of view. We define our set of parameters \mathcal{P} as:

$$\mathcal{P} = \{th_p, th_n, T_r, F_v\} \quad (1)$$

where:

- th_p represents the event threshold for positive events, *i.e.* an increase in logarithmic intensity, controlling the sensitivity of event generation.
- th_n represents the event threshold for negative occurrences, thus also adapting sensor sensitivity.
- T_r denotes the refractory period, affecting the temporal resolution of events.
- F_v stands for the field of view (FoV), determining the angular extent of the sensor coverage.

3.2.2 Configurations

We define a comprehensive set of experimental configurations $\mathcal{E} = \{\mathcal{E}_{base}, \mathcal{E}_1, \dots, \mathcal{E}_{13}\}$, where each configuration $\mathcal{E}_i \in \mathcal{E}$ represents a unique combination of event camera parameters. For any configuration \mathcal{E}_i and parameter $p \in \mathcal{P}$, we denote the value of p in \mathcal{E}_i as $p(\mathcal{E}_i)$.

We define a standard or default configuration, referred to as \mathcal{E}_{base} , which we use as the baseline to adapt the parameters in the subsequent configurations. This is the default sensor configuration provided by CARLA, except for the refractory period, which has been set to a very low value (0.01ms) instead of the default value of 0. In configurations $\mathcal{E}_i, i \in \{1, 2, \dots, 9\}$, we focus on varying a single setting $p \in \mathcal{P}$ over 3 distinct values for each, producing a distribution over each parameter space. At this stage,

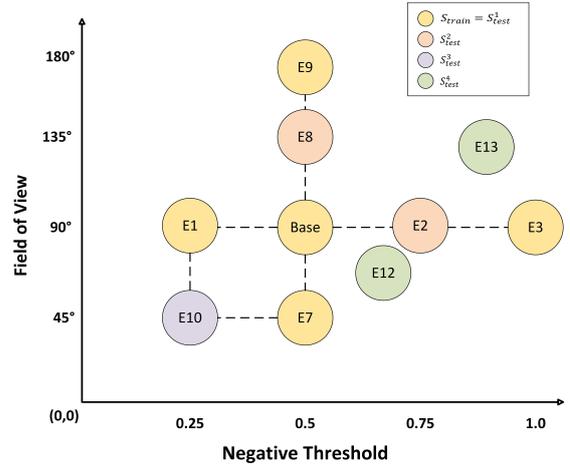


Figure 3: Simplified overview of the design and positioning of the distinct test sets in the parameter space of two parameters - the field of view and negative threshold. The training set and the first test set consist of distinct configurations varying along each dimension. The second test set consists of configurations obtained by interpolating along each dimension, whereas the third set is obtained by combining seen parameters into unique configurations. The fourth set represents the unseen parameter space for each parameter.

we do not consider $th_p(\mathcal{E}_i), th_n(\mathcal{E}_i)$ independently of each other, and vary them simultaneously. For settings $\mathcal{E}_i, i \in \{10, 11, 12, 13\}$, we fluctuate all the parameters simultaneously to generate unique distributions over the parameter space. Specifically for \mathcal{E}_{11} and \mathcal{E}_{13} we consider each of $th_p(\mathcal{E}_i), th_n(\mathcal{E}_i)$ independently, and set them at contrasting values. The specific parameter values for each configuration are detailed in Table 1 and visual examples for some of the configurations are displayed in Fig. 2.

3.3 Dataset Partitioning

Our experimental framework is designed to rigorously assess the generalization capabilities of event-based vision models across various event camera configurations. By leveraging the synthetic dataset generated in a controlled simulated environment, we partition our data into distinct training and test sets to design a systematic evaluation of multiple aspects of generalization. Fig. 3 provides a simplified visual overview of the main design idea.

1. Training Set:

This set, denoted by S_{train} , encompasses a diverse range of parameter combinations, serving as the foundation for model training and establishing the baseline performance. The idea is to use the lower and higher bounds of the range for each parameter $p_i \in \mathcal{P}$ for training to expose the model to the extreme variations

in the data distribution. This is supplemented by *base* configuration, \mathcal{E}_{base} , on top of which the variations are levied. We hypothesize that by learning the patterns of the limits of the data, the model can better adapt to unseen intermediate values. We compose this set as:

$$S_{train} = \{\mathcal{E}_{base}, \mathcal{E}_1, \mathcal{E}_3, \mathcal{E}_4, \mathcal{E}_6, \mathcal{E}_7, \mathcal{E}_9\} \quad (2)$$

2. Test Sets:

We define four test sets, each designed to evaluate a specific aspect of model generalization:

(a) Intra-distribution Generalization:

We use the same sensor characteristics for this test set as are used by our training set S_{train} . By employing the same parameter distribution, but a different split of the data (*e.g.* different towns), this approach assesses the model’s ability to generalize within the training parameter space. It is aimed at establishing a baseline performance and demonstrating the model’s robustness to external fluctuations, like geographical features, lighting conditions, and traffic density, within familiar non-divergent settings. Defined as:

$$S_{test}^1 = S_{train} \quad (3)$$

This train-test split is representative of what is commonly used in the literature, *i.e.*, you have data from a single source and partition it. Please note that \mathcal{E}_i describes the sensor configuration only; the training and test splits are generated for distinct towns.

(b) Single Parameter Perturbation:

In this setup, each configuration is deliberately designed to differ from all training configurations by exactly one parameter. This focused approach is aimed at isolating the influence each parameter has on the model’s performance by eliminating confounding factors. It is crucial for uncovering targeted insights into the model’s sensitivity, which can guide further model refinement. Furthermore, we can investigate if the model indeed learns to interpolate a single parameter between the provided limits during training. Formally:

$$S_{test}^2 = \{\mathcal{E}_2, \mathcal{E}_5, \mathcal{E}_8\} \quad (4)$$

(c) Distinct Configuration Derived from Observed Parameters:

Each parameter is sampled from the well-defined training distribution, but their combined state maps them to a novel resulting distribution. It aims to assess the model’s potential to interpolate between familiar configurations, demonstrating the model’s understanding of relationships between parameters, instead of overfitting to a singular parameter space. Evaluating its capacity to handle the complex sensor dynamics resulting from the interplay of multiple internal

factors is essential for the final use case, as each parameter can be adjusted independently by a controller, which would produce unique distributions in the combined space. Formally put:

$$S_{test}^3 = \{\mathcal{E}_{10}, \mathcal{E}_{11}\} \quad (5)$$

(d) Arbitrary Combinations of Unseen Parameters:

This test evaluates the model’s ability to perform beyond the known parameter values of the training data by employing entirely unseen individual parameters. It aims to assess the model’s capability and robustness when operating under significant shifts in the event camera’s operating characteristics. This is especially relevant for our ultimate objective of developing a dynamically adaptive event sensor, since it would theoretically be possible for the sensor to assume any permitted state. Ensuring the downstream model performs optimally, invariant to the sensor characteristics, is critical. Formally defined as:

$$S_{test}^4 = \{\mathcal{E}_{12}, \mathcal{E}_{13}\} \quad (6)$$

3.4 Network Architecture

We use two off-the-shelf event-based object detectors, Recurrent Vision Transformers (RVT) (Gehrig and Scaramuzza, 2023) and State Space Models (SSMs) (Zubic et al., 2024), for our experiments and evaluation. An overview of their concepts and architectures is discussed here.

3.4.1 Recurrent Vision Transformers (RVT)

The proposed architecture (Gehrig and Scaramuzza, 2023) employs a combination of convolutional layers, self-attention, and LSTMs to preserve and extract spatio-temporal information from dense tensor representations of the events.

Overlapping convolutional kernels are used to embed spatial relations between feature positions and for downsampling output maps of preceding layers. Self-attention is applied in two stages to cater to both local feature interactions using block attention and global feature interactions using grid attention. This strategy of global dilation of features counteracts the quadratic complexity that would have been otherwise required to apply self-attention over the entire feature map. Features across time-steps are aggregated with LSTMs initialized with states from the preceding time-steps.

This architecture block is repeated four times sequentially, and the hidden states of the LSTM, from the second to the fourth layer, are employed as features for the object detection head.

3.4.2 State Space Models (SSMs)

The key characteristic of (Zubic et al., 2024) involves replacing the LSTM blocks in the above-discussed architecture with State Space Model layers. This allows for faster parallel training utilizing convolutions and achieves a 33% more rapid training than employing LSTMs, by avoiding non-linearities which necessitated sequential processing of the data.

It also addresses the issue of performance degradation when employed at frequencies (temporal windows) divergent from the training frequency. Learnable timescale parameters are introduced, along with anti-aliasing strategies when the kernel bandwidth exceeds the Nyquist frequency.

Additionally, at higher sampling frequencies (beyond the Nyquist frequency), aliasing can severely impact performance. This is handled by a frequency-selective masking where the kernel focuses on lower frequencies, smoothening high-frequency components surpassing an empirical threshold. An additional H_2 norm regularization suppresses the frequency response of the system beyond a range given by $[w_{min}, w_{max}]$, during the computation of the norm of the system’s transfer function.

3.5 Implementation Details

We train and compare the performance of our joint training set against a base model for both the above architectures. All models are trained with a 32-bit precision over 400k steps, and employ the Adam optimizer and the OneCycle learning rate scheduler as in Gehrig and Scaramuzza (2023); Zubic et al. (2024). The best-performing checkpoint on the validation dataset is used for testing across all scenarios.

Random horizontal flips and zoom-in, zoom-out procedures are used to supplement the data. For the event-based representations, we utilize fixed temporal windows of 50ms, and discretize them over $n_{bins} = 10$ bins. We use a batch size of 8, a sequence length of 5, and a global learning rate of $2e^{-4}$ for all our experiments. The YOLOX (Ge et al., 2021) detection head incorporates an IOU loss, along with the regression loss, which is aggregated over batches and sequences in each step. The H_2 norm is additionally incorporated with the loss for the models using SSMs.

4 EXPERIMENTS AND RESULTS

We present and discuss our results in this subsection, providing a comprehensive analysis of the model’s performance across the previously designed test set-

tings. We employ standard evaluation metrics to quantify and compare all models, providing a leveled ground for further analysis.

4.1 Evaluation Framework

For each test set S_{test}^k , we compute a set of performance metrics $\mathcal{M} = \{m_1, m_2, \dots, m_n\}$. We use the COCO evaluation metrics (Lin et al., 2014) for object detection, where the Average Precision (AP) is the primarily employed metric. We define $\mathcal{M} = \{AP, AP_{50}, AP_{75}, AP_L, AP_M\}$. Here, AP is used as a shorthand for AP@[0.50:0.95], which reports the average over IoU thresholds from 0.50 to 0.95, and AP₅₀ and AP₇₅ are the AP scores for IoU values of 0.50 and 0.75, respectively. AP_M and AP_L refer to the scores for medium ($32 < \text{side length} < 96$ pixels) and large objects ($\text{side length} > 96$ pixels), respectively. We aim to provide an extensive knowledge of the model’s generalization capabilities, strengths, and limits in the context of event-based vision tasks across many event camera configurations by methodically investigating these performance indicators throughout our carefully prepared test sets.

The performance of a model f on S_{test}^k is defined:

$$\text{Score}_k(f) = \{\mathbb{E}_{\mathcal{E}_i \in S_{test}^k} [m_j(f, \mathcal{E}_i)] \mid m_j \in \mathcal{M}\} \quad (7)$$

where $m_j(f, \mathcal{E}_i)$ denotes the performance of model f on configuration \mathcal{E}_i according to metric m_j .

4.2 Sensor Generalization

We delve deeper into the outcomes in each condition, highlighting our observations, the strengths of our model, and areas that warrant further improvement. Fig. 4 presents a visual overview of the predictions of the models on different configurations.

Table 2 presents the AP scores obtained by the static-trained base model and our proposed expanded model for the RVT (Gehrig and Scaramuzza, 2023) and SSMS (Zubic et al., 2024) models for each of the configurations. We refer to the models trained only on \mathcal{E}_{base} as the static-trained model, and the performance of all four models (\mathcal{E}_{base} - and S_{train} -trained RVT-B and SSMS-B) on configuration \mathcal{E}_{base} as the baseline performance for each of them, respectively.

4.2.1 Performance across Thresholds

The threshold parameter, which governs the sensitivity of event generation, emerges as a critical factor influencing model robustness. When the threshold is set to a lower value, as in configuration \mathcal{E}_1 , the system generates a higher number of events. Notably, this

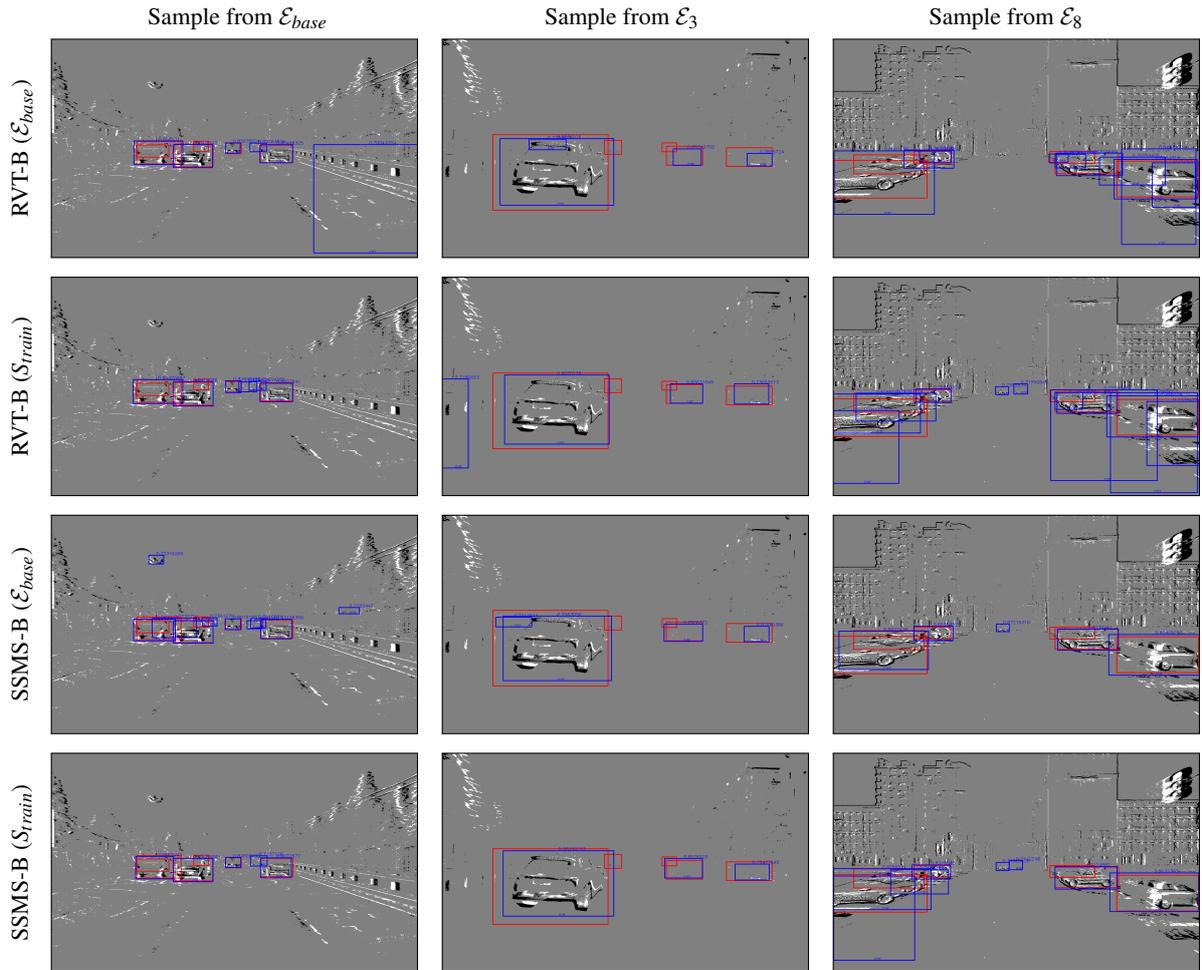


Figure 4: Comparison of results between the different models - RVT-B and SSMS-B trained on \mathcal{E}_{base} and S_{train} . The blue boxes indicate the predicted bounding boxes, and the red boxes represent the ground truth. Each of the three samples represents different sensor configurations to illustrate the performance across varying scenarios.

increase in event frequency does not adversely affect the performance of the static-trained model. In fact, for SSMs, the static-trained model even demonstrates improved results compared to its performance under the baseline threshold (\mathcal{E}_{base}). This counterintuitive outcome may be attributed to the model’s ability to leverage the richer temporal information provided by the denser event stream. Additionally, the event representation method employed plays a crucial role here. The Stacked Histogram representation with fixed temporal bins implements automatic saturation for high-density events through a clipping value.

Furthermore, our model exhibits a consistent, albeit modest, improvement over the static-trained baseline under these conditions *i.e.* $\sim 1\%$ for RVT and $\sim 2\%$ for SSMs. This suggests that the proposed training approach is capable of extracting additional value from the increased data density, even when the

static model already performs well.

However, the scenario changes dramatically as the threshold increases, resulting in sparser event data (as seen in \mathcal{E}_3). Here, the static-trained model’s performance deteriorates significantly, with a reduction of $\sim 23\%$ relative to the baseline. In contrast, our model, which was exposed to a broader range of event densities during training, demonstrates a more graceful degradation, *i.e.* only $\sim 15\%$ below baseline, thereby achieving an $\sim 8\%$ performance gain over the static model in this challenging regime. This finding underscores the importance of training with diverse event densities to enhance model resilience under sparse data conditions.

The robustness of our model is further highlighted in out-of-distribution scenarios, such as configuration \mathcal{E}_2 , where the threshold is set to an unseen value. Even in this setting, the model maintains its performance

Table 2: Comparison of performance (AP scores in %) for each configuration.

| Train/Test | | RVT-B | | SSMS-B | |
|--------------|----------------------|----------------------|--------------|----------------------|--------------|
| Config. | | \mathcal{E}_{base} | S_{train} | \mathcal{E}_{base} | S_{train} |
| S_{test}^1 | \mathcal{E}_{base} | 45.63 | 44.12 | 49.10 | 50.58 |
| | \mathcal{E}_1 | 44.49 | 45.94 | 51.99 | 53.90 |
| | \mathcal{E}_3 | 23.74 | 30.23 | 26.72 | 33.01 |
| | \mathcal{E}_4 | 45.13 | 44.72 | 48.97 | 50.55 |
| | \mathcal{E}_6 | 44.61 | 44.13 | 48.73 | 50.05 |
| | \mathcal{E}_7 | 31.49 | 35.61 | 34.68 | 41.31 |
| | \mathcal{E}_9 | 7.32 | 17.10 | 7.08 | 17.58 |
| | Avg. | 34.63 | 37.41 | 38.18 | 42.42 |
| S_{test}^2 | \mathcal{E}_2 | 35.69 | 37.69 | 38.69 | 42.03 |
| | \mathcal{E}_5 | 45.56 | 44.48 | 48.97 | 50.98 |
| | \mathcal{E}_8 | 31.21 | 35.95 | 32.09 | 41.27 |
| | Avg. | 37.49 | 39.37 | 39.92 | 44.76 |
| S_{test}^3 | \mathcal{E}_{10} | 33.92 | 40.00 | 40.68 | 48.14 |
| | \mathcal{E}_{11} | 34.22 | 37.70 | 40.54 | 43.13 |
| | Avg. | 34.07 | 38.85 | 40.61 | 45.64 |
| S_{test}^4 | \mathcal{E}_{12} | 34.08 | 37.30 | 35.79 | 40.86 |
| | \mathcal{E}_{13} | 27.09 | 31.03 | 30.88 | 37.26 |
| | Avg. | 30.58 | 34.16 | 33.33 | 39.06 |

advantage, outperforming the static baseline by $\sim 2\%$ in RVT and $\sim 4\%$ in SSMS. This result attests to the generalization capabilities conferred by the employed training methodology.

4.2.2 Invariance to Refractory Period

Varying the refractory period does not have a significant effect on model performance across all tested configurations. The static-trained model maintains stable results, indicating that its learned representations are not particularly sensitive to this parameter. For the expanded model, the effect is similarly negligible, with only a slight decrease in RVT performance and a modest $\sim 1 - 2\%$ improvement in SSMS. Importantly, this stability persists even when the values fall outside the training distribution, suggesting that both models possess a degree of invariance to this parameter that is desirable for practical deployment in variable hardware environments.

4.2.3 Performance across Fields of View

The field-of-view (FoV) parameter, which determines the angular extent of the sensor’s perceptual field, has a pronounced impact on detection performance. Altering the FoV fundamentally changes the geometry of the observed scene: increasing the FoV stretches

and flattens object appearances for objects, and performs an indirect scaling operation - an increased field of view makes an object appear distant and therefore smaller. These perspective distortions can place test samples significantly out of the training distribution, challenging the model’s generalization ability.

When the FoV is narrowed to 45° (as in \mathcal{E}_7), the static-trained model’s performance drops precipitously by $\sim 14\%$ compared to the baseline. In contrast, our expanded model exhibits a more moderate decline of $\sim 9\%$, and the SSMS variant achieves a relative gain of $\sim 7\%$ over the static model. This resilience is likely attributable to the model’s exposure to a wider range of FoV settings during training, which enables it to accommodate geometric distortions better.

At the other extreme, widening the FoV to 160° in configuration \mathcal{E}_9 leads to the most severe performance degradation across all settings. Nevertheless, when the training distribution includes such wide FoV examples, our model can outperform the static baseline by as much as $\sim 10\%$. Even in intermediate, previously unseen FoV settings (e.g., 135° in \mathcal{E}_8), our model maintains a substantial performance margin, achieving up to $\sim 9\%$ improvement in SSMS. These results collectively underscore the importance of training with a diverse range of geometric configurations to ensure robust performance.

4.2.4 Performance on Out-of-Distribution Configurations

When tested on joint parameter combinations (\mathcal{E}_{10}) where individual parameters resided within the training distribution but their inter-dependencies deviated from training patterns, our methodology exhibited moderate performance degradation: a $\sim 4\%$ reduction for RVT and $\sim 2\%$ for SSMS relative to baseline. This contrasts sharply with the static-trained model, which suffered disproportionately larger declines under equivalent conditions. Notably, the performance gap widens when both individual parameters and their joint configurations fall outside the training domain, as in \mathcal{E}_{12} , with our approach maintaining a consistent $\sim 4 - 6\%$ advantage across metrics. These results suggest that while parameter interdependence introduces complexity, the proposed training strategy mitigates cascading errors through learned representations that decouple spurious parameter correlations.

4.2.5 Response to Asymmetric Positive and Negative Thresholds

A critical stress test evaluated model robustness to asymmetric event-generation thresholds, where posi-

Table 3: Complete Performance Comparison on the RVT-B Model and the SSMS-B Model. We observe that the SSMS-B Model trained on S_{train} performs the best uniformly in each score metric, across all test sets.

| Score _k (<i>f</i>) | AP(<i>m</i> ₁) | | AP ₅₀ (<i>m</i> ₂) | | AP ₇₅ (<i>m</i> ₃) | | AP _L (<i>m</i> ₄) | | AP _M (<i>m</i> ₅) | |
|--|-----------------------------|------------------------------|--|------------------------------|--|------------------------------|---|------------------------------|---|------------------------------|
| | RVT- B | SSMS- B | RVT- B | SSMS- B | RVT- B | SSMS- B | RVT- B | SSMS- B | RVT- B | SSMS- B |
| Score ₁ (\mathcal{E}_{base}) | 34.6 ± 14.7 | 38.2 ± 16.6 | 59.7 ± 17.9 | 60.9 ± 19.8 | 37.4 ± 19.7 | 42.8 ± 21.6 | 44.0 ± 18.0 | 49.6 ± 20.4 | 30.0 ± 13.4 | 32.0 ± 14.9 |
| Score ₁ (S_{train}) | 37.4 ± 10.7 | 42.4 ± 13.1 | 63.4 ± 11.3 | 65.3 ± 14.2 | 40.7 ± 15.5 | 47.7 ± 17.5 | 47.4 ± 14.1 | 53.8 ± 17.4 | 32.3 ± 9.3 | 36.4 ± 11.4 |
| Score ₂ (\mathcal{E}_{base}) | 37.5 ± 7.3 | 39.9 ± 8.5 | 64.6 ± 6.3 | 65.3 ± 6.9 | 39.7 ± 12.4 | 44.0 ± 12.1 | 46.3 ± 12.8 | 46.9 ± 16.1 | 33.5 ± 6.0 | 35.9 ± 6.0 |
| Score ₂ (S_{train}) | 39.4 ± 4.5 | 44.8 ± 5.4 | 66.7 ± 3.6 | 68.9 ± 4.4 | 43.1 ± 7.8 | 51.3 ± 7.1 | 47.4 ± 10.7 | 54.4 ± 12.3 | 35.6 ± 4.6 | 39.8 ± 5.6 |
| Score ₃ (\mathcal{E}_{base}) | 34.1 ± 0.2 | 40.6 ± 0.1 | 62.3 ± 4.5 | 66.7 ± 2.7 | 34.4 ± 3.7 | 45.2 ± 0.0 | 39.3 ± 3.3 | 54.9 ± 4.4 | 27.6 ± 1.5 | 31.9 ± 0.6 |
| Score ₃ (S_{train}) | 38.9 ± 1.6 | 45.6 ± 3.5 | 64.9 ± 2.7 | 69.4 ± 3.9 | 43.4 ± 2.3 | 51.9 ± 5.1 | 52.3 ± 0.6 | 60.9 ± 0.8 | 30.7 ± 0.6 | 36.0 ± 2.4 |
| Score ₄ (\mathcal{E}_{base}) | 30.6 ± 4.9 | 33.3 ± 3.5 | 57.5 ± 2.6 | 59.9 ± 0.4 | 29.7 ± 9.2 | 34.2 ± 7.1 | 40.8 ± 14.2 | 43.2 ± 16.3 | 24.5 ± 0.6 | 28.1 ± 4.2 |
| Score ₄ (S_{train}) | 34.2 ± 4.4 | 39.1 ± 2.5 | 62.5 ± 0.0 | 64.9 ± 0.6 | 34.5 ± 9.4 | 41.9 ± 4.9 | 43.7 ± 15.3 | 49.0 ± 14.3 | 28.9 ± 1.5 | 33.5 ± 4.1 |

tive and negative polarity triggers operated under distinct criteria (\mathcal{E}_{11} , \mathcal{E}_{13}), which was completely absent from all training data. The static-trained model exhibited pronounced sensitivity, with performance plummeting by $\sim 11\%$ (RVT) and $\sim 9\%$ (SSMs) when tested on threshold pairs independently overlapping its training distribution. Under fully novel asymmetries, degradation intensified to 18–20% relative to baseline. In contrast, our expanded model demonstrated superior adaptability, sustaining only $\sim 7\%$ (RVT) initial degradation, which increased to 13–14% under extreme distributional mismatch. The systematic 4–5% performance gap across all asymmetric configurations underscores the method’s capacity to generalize beyond symmetric threshold paradigms, a capability likely rooted in its exposure to diverse event-density variations during training, which fosters robust sensor-parameter invariance.

4.2.6 Differences between RVT and SSMS

In general, we observe that the SSMS model, paired with our training strategy, performs the best across all sensor scenarios and learns robust parameter disentanglement. It has a greater capability to generalize between the representations of the vast data distribution, compared to the RVT.

Notably, even the baseline performance (on \mathcal{E}_{base}) of the static-trained model improves by $\sim 4\%$ when employing SSMS instead of RVT. For the same

model, testing on \mathcal{E}_1 , which has a significantly denser event representation due to its low threshold, this performance gap increases to $\sim 7\%$, which our model further slightly widens. This shows that the frequency-generalizing adaptations adopted in SSMS also extend their profitability to the handling of varying event densities. However, when the data is significantly out-of-distribution, as in \mathcal{E}_9 , any such significant performance gap between the models vanishes, both for the baseline and our expanded model.

4.2.7 Other Metrics

We summarily present the results for all metrics, $m \in \mathcal{M}$, for each of the test sets, in Table 3. We present the collective scores $\text{Score}_k(f)$, instead of a detailed breakdown, as all individual metrics follow a similar pattern to the Average Precision (AP) scores, outlined and discussed at length.

5 CONCLUSION

In this work, we present a systematic approach towards sensor-invariant event-based object detection, addressing the fundamental challenges posed by varying sensor configurations. We accumulate an expansive simulated dataset to study the impact of varying sensor characteristics on the performance in downstream tasks. We highlight the limitations and weak-

nesses of a static-trained model in adapting to changes in signal distribution. By leveraging multi-source domain generalization, we demonstrated that training across diverse sensor parameters enhances the adaptability of event-based vision models, mitigating performance degradation due to fluctuating sensor characteristics. To the best of our knowledge, this is the first work to carry out such an expansive study into underlying sensor characteristics and the expansion of domain generalization into this scope of event data.

Through rigorous evaluation across a range of carefully designed test settings, we established the effectiveness of our employed strategy in expanding model robustness beyond standard configurations. Our findings underscore the significance of dataset diversity in training sensor-agnostic models, thereby paving the way for more reliable and scalable event-based perception systems. At the same time, we analyze failure points of our model, which warrant further research and improvement.

We believe that this work paves the way for effective future research focusing on developing dynamic adaptation mechanisms that allow for real-time tuning of sensor configurations based on environmental factors. Additionally, incorporating cross-modal fusion with RGB-based systems may further enhance generalization and improve event-camera integration into mainstream applications.

ACKNOWLEDGEMENTS

This work is a part of the project *SNACE*, which is partially funded by the Volkswagen Foundation under the profile area *Exploration* and the funding framework *NEXT: Neuromorphic Computing*.

REFERENCES

- Baldwin, R. W., Liu, R., Almatrafi, M., Asari, V., and Hirakawa, K. (2022). Time-ordered recent event (tore) volumes for event cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):2519–2532.
- Bendig, K., Schuster, R., and Stricker, D. (2023). On the future of training spiking neural networks. In *ICPRAM*, pages 466–473.
- Bi, Y., Chadha, A., Abbas, A., Boursoulatze, E., and Andreopoulos, Y. (2019). Graph-based object classification for neuromorphic vision sensing. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 491–501.
- Bu, T., Ding, J., Yu, Z., and Huang, T. (2022). Optimized potential initialization for low-latency spiking neural networks. In *Proceedings of the AAAI conference on artificial intelligence*.
- Cordone, L., Miramond, B., and Thierion, P. (2022). Object detection with spiking neural networks on automotive event data. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- De Tournemire, P., Nitti, D., Perot, E., Migliore, D., and Sironi, A. (2020). A large scale event-based detection dataset for automotive. *arXiv preprint arXiv:2001.08499*.
- Deng, S. and Gu, S. (2021). Optimal conversion of conventional artificial neural networks to spiking neural networks. *arXiv preprint arXiv:2103.00476*.
- Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., and Koltun, V. (2017). Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR.
- Eckmann, S., Klimmasch, L., Shi, B. E., and Triesch, J. (2020). Active efficient coding explains the development of binocular vision and its failure in amblyopia. *Proceedings of the National Academy of Sciences*, 117(11):6156–6162.
- Fan, L., Li, Y., Shen, H., Li, J., and Hu, D. (2024a). From dense to sparse: low-latency and speed-robust event-based object detection. *IEEE Transactions on Intelligent Vehicles*.
- Fan, Y., Zhang, W., Liu, C., Li, M., and Lu, W. (2024b). Sfod: Spiking fusion object detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17191–17200.
- Ge, Z., Liu, S., Wang, F., Li, Z., and Sun, J. (2021). YoloX: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*.
- Gehrig, D., Gehrig, M., Hidalgo-Carri6, J., and Scaramuzza, D. (2020). Video to events: Recycling video datasets for event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3586–3595.
- Gehrig, D. and Scaramuzza, D. (2022). Pushing the limits of asynchronous graph-based object detection with event cameras. *arXiv preprint arXiv:2211.12324*.
- Gehrig, M. and Scaramuzza, D. (2023). Recurrent vision transformers for object detection with event cameras. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13884–13893.
- Guo, Y., Huang, X., and Ma, Z. (2023). Direct learning-based deep spiking neural networks: a review. *Frontiers in Neuroscience*, 17:1209795.
- Hamaguchi, R., Furukawa, Y., Onishi, M., and Sakurada, K. (2023). Hierarchical neural memory network for low latency event processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22867–22876.
- Jian, D. and Rostami, M. (2023). Unsupervised domain adaptation for training event-based networks using contrastive learning and uncorrelated conditioning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 18721–18731.

- Kim, S., Park, S., Na, B., and Yoon, S. (2020). Spiking-yolo: spiking neural network for energy-efficient object detection. In *Proceedings of the AAAI conference on artificial intelligence*, pages 11270–11277.
- Li, J., Li, J., Zhu, L., Xiang, X., Huang, T., and Tian, Y. (2022). Asynchronous spatio-temporal memory network for continuous event-based object detection. *IEEE Transactions on Image Processing*, 31:2975–2987.
- Li, Y. and Zeng, Y. (2022). Efficient and accurate conversion of spiking neural network with burst spikes. *arXiv preprint arXiv:2204.13271*.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. (2014). Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer.
- Lu, D., Kong, L., Lee, G. H., Chane, C. S., and Ooi, W. T. (2024). Flexevent: Event camera object detection at arbitrary frequencies. *arXiv preprint arXiv:2412.06708*.
- Maqueda, A. I., Loquercio, A., Gallego, G., García, N., and Scaramuzza, D. (2018). Event-based vision meets deep learning on steering prediction for self-driving cars. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5419–5427.
- Nam, Y., Mostafavi, M., Yoon, K.-J., and Choi, J. (2022). Stereo depth from events cameras: Concentrate and focus on the future. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6114–6123.
- Peng, Y., Zhang, Y., Xiong, Z., Sun, X., and Wu, F. (2023). Get: Group event transformer for event-based vision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6038–6048.
- Perot, E., De Tournemire, P., Nitti, D., Masci, J., and Sironi, A. (2020). Learning to detect objects with a 1 megapixel event camera. *Advances in Neural Information Processing Systems*, 33:16639–16652.
- Planamente, M., Plizzari, C., Cannici, M., Ciccone, M., Strada, F., Bottino, A., Matteucci, M., and Caputo, B. (2021). Da4event: towards bridging the sim-to-real gap for event cameras using domain adaptation. *IEEE Robotics and Automation Letters*, 6(4):6616–6623.
- Rebecq, H., Horstschaefer, T., and Scaramuzza, D. (2017). Real-time visual-inertial odometry for event cameras using keyframe-based nonlinear optimization. In *British Machine Vision Conference (BMVC)*.
- Schaefer, S., Gehrig, D., and Scaramuzza, D. (2022). Aegnn: Asynchronous event-based graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12371–12381.
- Silva, D. A., Smagulova, K., Elsheikh, A., Fouda, M. E., and Eltawil, A. M. (2025). A recurrent yolov8-based framework for event-based object detection. *Frontiers in Neuroscience*, 18:1477979.
- Sironi, A., Brambilla, M., Bourdis, N., Lagorce, X., and Benosman, R. (2018). Hats: Histograms of averaged time surfaces for robust event-based object classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1731–1740.
- Sprague, W. W., Cooper, E. A., Tošić, I., and Banks, M. S. (2015). Stereopsis is adaptive for the natural environment. *Science advances*, 1(4):e1400254.
- Stoffregen, T., Scheerlinck, C., Scaramuzza, D., Drummond, T., Barnes, N., Kleeman, L., and Mahony, R. (2020). Reducing the sim-to-real gap for event cameras. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVII 16*, pages 534–549. Springer.
- Su, Q., Chou, Y., Hu, Y., Li, J., Mei, S., Zhang, Z., and Li, G. (2023). Deep directly-trained spiking neural networks for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6555–6565.
- Sundar, V., Dutson, M., Ardelean, A., Bruschini, C., Charbon, E., and Gupta, M. (2024). Generalized event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25007–25017.
- Wang, J., Lan, C., Liu, C., Ouyang, Y., Qin, T., Lu, W., Chen, Y., Zeng, W., and Yu, P. S. (2022). Generalizing to unseen domains: A survey on domain generalization. *IEEE transactions on knowledge and data engineering*, 35(8):8052–8072.
- Wang, Y.-K., Wang, S.-E., and Wu, P.-H. (2023). Spike-event object detection for neuromorphic vision. *IEEE Access*, 11:5215–5230.
- Yao, M., Gao, H., Zhao, G., Wang, D., Lin, Y., Yang, Z., and Li, G. (2021). Temporal-wise attention spiking neural networks for event streams classification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10221–10230.
- Zhang, H., Wang, X., Xu, C., Wang, X., Xu, F., Yu, H., Yu, L., and Yang, W. (2024). Frequency-adaptive low-latency object detection using events and frames. *arXiv preprint arXiv:2412.04149*.
- Zhang, J., Dong, B., Zhang, H., Ding, J., Heide, F., Yin, B., and Yang, X. (2022). Spiking transformers for event-based single object tracking. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 8801–8810.
- Zhao, J., Zhang, S., and Huang, T. (2022). Transformer-based domain adaptation for event data classification. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4673–4677. IEEE.
- Zhou, K., Liu, Z., Qiao, Y., Xiang, T., and Loy, C. C. (2022). Domain generalization: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4396–4415.
- Zubic, N., Gehrig, M., and Scaramuzza, D. (2024). State space models for event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5819–5828.