



## OPEN ACCESS

## EDITED BY

Akshay Agarwal,  
Indian Institute of Science Education and  
Research, Bhopal, India

## REVIEWED BY

Chao Qi,  
Jiangsu Academy of Agricultural Sciences  
(JAAS), China  
Maria Lombardi,  
Italian Institute of Technology (IIT), Italy

## \*CORRESPONDENCE

Mohammed E. Fouda  
✉ foudam@uci.edu

RECEIVED 08 August 2024

ACCEPTED 20 December 2024

PUBLISHED 22 January 2025

## CITATION

Silva DA, Smagulova K, Elsheikh A, Fouda ME  
and Eltawil AM (2025) A recurrent  
YOLOv8-based framework for event-based  
object detection.  
*Front. Neurosci.* 18:1477979.  
doi: 10.3389/fnins.2024.1477979

## COPYRIGHT

© 2025 Silva, Smagulova, Elsheikh, Fouda and  
Eltawil. This is an open-access article  
distributed under the terms of the [Creative  
Commons Attribution License \(CC BY\)](#). The  
use, distribution or reproduction in other  
forums is permitted, provided the original  
author(s) and the copyright owner(s) are  
credited and that the original publication in  
this journal is cited, in accordance with  
accepted academic practice. No use,  
distribution or reproduction is permitted  
which does not comply with these terms.

# A recurrent YOLOv8-based framework for event-based object detection

Diego A. Silva<sup>1</sup>, Kamilya Smagulova<sup>1</sup>, Ahmed Elsheikh<sup>2</sup>,  
Mohammed E. Fouda<sup>3\*</sup> and Ahmed M. Eltawil<sup>1</sup>

<sup>1</sup>Communication and Computing Systems Lab, Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia, <sup>2</sup>Mathematics and Engineering Physics Department, Faculty of Engineering, Cairo University, Giza, Egypt, <sup>3</sup>Compumacy for Artificial Intelligence Solutions, Cairo, Egypt

Object detection plays a crucial role in various cutting-edge applications, such as autonomous vehicles and advanced robotics systems, primarily relying on conventional frame-based RGB sensors. However, these sensors face challenges such as motion blur and poor performance under extreme lighting conditions. Novel event-based cameras, inspired by biological vision systems, offer a promising solution with superior performance in fast-motion and challenging lighting environments while consuming less power. This work explores the integration of event-based cameras with advanced object detection frameworks, introducing Recurrent YOLOv8 (ReYOLOv8), a refined object detection framework that enhances a leading frame-based YOLO detection system with spatiotemporal modeling capabilities by adding recurrency. ReYOLOv8 incorporates a low-latency, memory-efficient method for encoding event data called Volume of Ternary Event Images (VTEI) and introduces a novel data augmentation technique based on Random Polarity Suppression (RPS) optimized for event-based sensors and tailored to leverage the unique attributes of event data. The framework was evaluated using two comprehensive event-based datasets Prophesee's Generation 1 (GEN1) and Person Detection for Robotics (PEDRo). On the GEN1 dataset, ReYOLOv8 achieved mAP improvements of 5%, 2.8%, and 2.5% across nano, small, and medium scales, respectively, while reducing trainable parameters by 4.43% on average and maintaining real-time processing speeds between 9.2 ms and 15.5 ms. For the PEDRo dataset, ReYOLOv8 demonstrated mAP improvements ranging from 9% to 18%, with models reduced in size by factors of 14.5× and 3.8× and an average speed improvement of 1.67×. The results demonstrate the significant potential of bio-inspired event-based vision sensors when combined with advanced object detection frameworks. In particular, the ReYOLOv8 system effectively bridges the gap between biological principles of vision and artificial intelligence, enabling robust and efficient visual processing in dynamic and complex environments. The codes are available on GitHub at the following link <https://github.com/silvada95/ReYOLOv8>.

## KEYWORDS

object detection, YOLO, event-based cameras, data augmentation, autonomous driving

# 1 Introduction

Object Detection involves the dual processes of locating and categorizing objects within an image, serving as a critical function in a multitude of fields, including Autonomous Driving (Michaelis et al., 2019), Robotics (Xu et al., 2022), and Surveillance (Jha et al., 2021). Deep Learning algorithms primarily drive advancements in this area, with the You Only Look Once (YOLO) detectors, based on Convolutional Neural Networks (CNN), emerging as a prominent choice in academia and industry. Renowned for its real-time capabilities, YOLO stands out for its efficient performance with minimal parameters, as documented in various studies (Liu et al., 2020). Over time, YOLO has undergone iterations and enhancements, making it faster and more robust in handling object detection tasks (Terven et al., 2023).

In contemporary computer vision applications, the standard practice involves processing images captured by cameras that detect light in red, green, and blue wavelengths (RGB). While modern RGB sensors excel in providing high-resolution and detailed frames, they are susceptible to motion blur during high-speed movements, and their limited High Dynamic Range (HDR) poses challenges in complex lighting scenarios (Chen et al., 2020). In contrast, event-based cameras operate based on changes in illumination rather than absolute light levels, drawing inspiration from the human eye's visual data processing mechanism. This novel approach results in sparse data sequences comprising spatial coordinates, timestamps, and polarity information triggered by variations in light stimuli at specific pixels (Posch et al., 2014). Event-based cameras offer distinctive advantages, including ultra-low latency in the microsecond range, an HDR typically exceeding 120 dB, and power consumption in the milliwatt range. These characteristics render event-based cameras particularly suitable for time-critical tasks and challenging lighting conditions, making them a preferred choice in applications where swift responsiveness and adaptability to varying light environments are paramount (Gallego et al., 2022). However, processing event-based data requires development of novel techniques for feature learning and extraction (Lagorce et al., 2015; Annamalai et al., 2021).

The majority of existing object detectors tailored for event-based data primarily target autonomous driving scenarios, typically relying on datasets like Prophesee's Generation 1 (GEN1) (De Tournemire et al., 2020) and 1 MegaPixel (Perot et al., 2020), alongside robotics applications supported by the recently introduced Person Detection in Robotics (PEDRo) dataset (Boretti et al., 2023). These detectors are commonly categorized into two groups based on their approach to event stream handling: direct processing of sparse event streams and densification before processing. The former group includes Graph Neural Networks (GNNs) (Bi et al., 2019), Spiking Neural Networks (SNNs) (Cordone et al., 2022), and sparse CNNs (Messikommer et al., 2020). On the other hand, detectors in the second group first densify events before applying feature extractors such as CNNs (Perot et al., 2020) and transformers (Gehrig and Scaramuzza, 2023) and occasionally incorporate Recurrent Neural Networks (RNNs) or State Space Models (SSM) (Zubic et al., 2024) for modeling temporal relationships. On both categories, the feature extractors are combined with detection heads commonly found in

frame-based detection models, such as the YOLO family (Gehrig and Scaramuzza, 2023; Peng et al., 2023b; Zubić et al., 2023), RetinaNet (Lin et al., 2020), and Single-Shot Detector (SSD) (Liu et al., 2016), which are proven to provide good detection capability in event-based scenarios.

Despite recent advancements in GNNs and SNNs, detectors relying on densified event representations consistently outperform them by a significant margin, as evidenced in studies such as those by Gehrig and Scaramuzza (2024), Wang et al. (2024), and Zubić et al. (2023). Notably, top-performing detectors often leverage state-of-the-art detection heads borrowed from frame-based literature, leading to superior performance. Considering this trend, this research has opted to build upon the YOLOv8 framework (Jocher et al., 2020) as the foundation for the event-based detector development. YOLOv8's exceptional performance and real-time processing capabilities make it a compelling choice over other alternatives, with a focus on achieving enhanced performance with reduced parameters. Recognizing the success of integrating spatial feature extractors and RNNs in event data processing, as demonstrated in works like Perot et al. (2020), Li et al. (2022a), and Gehrig and Scaramuzza (2023), the YOLOv8 framework has been enhanced to incorporate Convolutional Long-Short Term Memory (ConvLSTM) cells, along with the implementation of Truncated-Backpropagation Through Time (T-BPTT) for RNN training (Shi et al., 2015).

A critical aspect of current event-based object detectors lies in the dense encodings used for effectively feeding event stream data into neural networks. Various encoding strategies have been proposed in the literature, each aiming to retain crucial information from event streams. These strategies range from simple projections on 2D planes based on event counting (Rebecq et al., 2017; Maqueda et al., 2018), timestamp manipulation techniques (Benosman et al., 2013; Sironi et al., 2018; Baldwin et al., 2019), to hybrid approaches combining both methods (Zhu et al., 2018a; El Shair et al., 2023). Other methodologies involve segmenting streams into spatiotemporal volumes (Zhu et al., 2018b; Gehrig and Scaramuzza, 2023; Nam et al., 2022; Peng et al., 2023a; Fan et al., 2024a), while learning-based encodings (Gehrig et al., 2019; Wang et al., 2023a), Bayesian optimization (Zubić et al., 2023), and the utilization of First-In-First-Out (FIFO) buffers (Baldwin et al., 2022; Liu et al., 2023) have also been proposed. Each of these approaches presents distinct trade-offs, impacting detection performance, encoding latency, and memory requirements associated with event inputs. While existing literature assesses the impact of these encoding choices on detection performance, inference time, and event processing rates, a comprehensive system-level perspective is lacking. To address this gap and introduce a memory-efficient fast encoding alternative, this work introduces a novel event representation titled Volume of Ternary Event Images (VTEI). A comparative analysis against closely related alternatives was conducted, not only focusing on parameters like latency but also assessing hardware-related factors such as data unit encoding size, memory footprint under various event rates, compression ratios, and bandwidth requirements.

Moreover, recognizing the lack of techniques that specifically address event-based feature augmentation, this study introduces a novel approach known as Random Polarity Suppression (RPS).

This method involves randomly suppressing all events associated with a particular polarity, enabling the detector to learn object-relevant features in a polarity-agnostic manner and mitigating potential biases in polarity distribution that could exist within the training dataset.

In summary, the contributions of this work are:

- A Recurrent-Convolutional Event-Based Object Detection network was introduced by means of the modification of the well-acknowledged real-time detector YOLOv8. The resulting framework, called Recurrent YOLOv8 (ReYOLOv8), was based on the addition of recurrent blocks and training with T-BPTT to the original framework, turning it capable of performing long-range spatiotemporal modeling;
- A fast and lightweight memory encoding called Volume of Ternary Event Images (VTEI) was proposed. This format is capable of retaining temporal information from event streams while presenting low latency, low bandwidth, high sparsity, and high compression ratio;
- A novel data augmentation technique based on Random Polarity Suppression (RPS) was introduced, showing success in improving the performance of the detection systems;
- The aforementioned contributions were merged into a single system, and validation of the resulting framework was performed across three different model scales over two real-world large-scale datasets. State-of-the-art performance for similar scale models was reported.

These contributions collectively advance the field of event-based neural processing, offering both theoretical and practical insights for object detection tasks in dynamic environments. The paper is organized as follows: Section 2 presents a review of the related works on event representations, detectors, and data augmentation techniques. Then, in Section 3, a discussion about the ideas proposed in this paper is performed. After that, the results are exposed in Section 4. Finally, Section 5 summarizes the main achievements of this work and provides some insights about future works.

## 2 Related works

### 2.1 Event representations

One of the most common and intuitive methods for event representation involves projecting events onto a 2D-pixel grid for modern CNN compatibility. An effective approach involves generating 2D grids based on timestamps (Benosman et al., 2013; Sironi et al., 2018; Baldwin et al., 2019). Along the lines of this 2D concept, Event Frames—or Histograms—rely on event counts at each pixel location (Rebecq et al., 2017), presenting in some cases channels separated by polarity (Maqueda et al., 2018) or combinations of channels integrating polarity and timestamp features (Zhu et al., 2018a; El Shair et al., 2023).

Preserving temporal information from events often involves constructing dense representations segmented into distinct temporal windows, subsequently stacked to construct a 3D tensor.

Voxel grids bin events across the time dimension, utilizing a bilinear kernel and interval normalization to weigh polarity contributions (Zhu et al., 2018b). In contrast, Stacked Histograms replace this kernel by a simple event counting (Gehrig and Scaramuzza, 2023). Mixed-Density Event Stacks (MDES) offer a variation where bins encode different event densities within a single window segment (Nam et al., 2022). Hyper Histograms split temporal windows into smaller units, creating channels based on polarity and timestamp histograms (Peng et al., 2023a). Event Temporal Images map events within the 0–255 range to create grayscale images, incorporating distinct ranges to capture differing positive and negative event distributions (Fan et al., 2024a).

Event Spike Tensor (EST) proposed an end-to-end learning process where MLPs are trained to find the best encoding according to a generalized 4D tensor, defined over the polarity and spatiotemporal domains (Gehrig et al., 2019). Asynchronous Attention Embedding employs an attention layer on events followed by a dilation-based 1D convolution used for data encoding (Li et al., 2022a). EventPillars, inspired by PointPillars, is a trainable representation that treats events similarly to point clouds (Wang et al., 2023a). Event Representation through Gromov-Wasserstein Optimization (ERGO) employs Bayesian optimization over categorical variables, leveraging the Gromov-Wasserstein Discrepancy (GWD) as a key metric to assess the effectiveness of a particular event representation (Zubić et al., 2023).

Time-ordered Recent Events (TORE) volumes utilize First-In-First-Out (FIFO) queues with a depth of  $K$ , establishing a direct mapping to each pixel at every polarity. These volumes can be generated asynchronously, without a predefined time window (Baldwin et al., 2022). Temporal Active Focus (TAF) aligns with TORE principles but integrates adaptive features for varying rates (Liu et al., 2023).

Given the asynchronous nature of events, a proposed approach involves encoding them as nodes within graphs, with the connections between nodes defined as edges. This methodology allows efficient processing using Graph Neural Networks (GNNs) (Bi et al., 2019). Voxel Cube introduced an alternative to event volumes where event accumulation within each micro-bin is binary, aiming to enhance temporal resolution specifically tailored for Spiking Neural Networks (SNNs) (Cordone et al., 2022). Group Tokens were specifically crafted for integration within Transformer-based architectures, involving the discretization of the event stream into intervals that are subsequently converted into patches. A 3 *times* 3 group convolution is then employed to embed the information into tokens effectively within this framework (Peng et al., 2023b).

In this work, a memory-efficient and rapid event representation called VTEI is proposed to contribute to the design of an efficient and lightweight object detection framework. VTEI leverages a spatiotemporal volume to preserve temporal information, similar to Voxel Grids and Stacked Histograms. However, unlike these methods, VTEI represents each data unit within the volume using a limited number of values, similar to MDES. This approach results in a final representation characterized by high sparsity, low memory usage, low bandwidth, and low latency. Furthermore, VTEI effectively preserves sub-temporal dynamics within a given time window using minimal polarity information.

## 2.2 Event-based object detectors

One of the pioneering works on event-based object detection, Asynet, proposed leveraging the intrinsic spatial sparsity of event data by converting synchronous networks to asynchronous ones (Messikommer et al., 2020). Recently, a Graph Neural Network (GNN) approach called Asynchronous Event-Based GNN (AEGNN) was introduced, modeling events as spatio-temporal graphs with events as nodes and connections between neighboring events as edges. Processing is conducted through graph pooling and graph convolutions (Schaefer et al., 2022). The potential of this approach has inspired the proposal of other similar networks (Sun and Ji, 2023; Gehrig and Scaramuzza, 2024).

Spiking Neural Networks (SNNs), driven by spikes analogous to events, are acknowledged for their low power consumption, making them suitable for event-based camera applications. A hybrid SNN-ANN architecture was proposed, utilizing end-to-end training to leverage event information without intermediate representations (Kugele et al., 2021). The first SNN validated on real-world event data incorporated spiking variants of VGG (Simonyan and Zisserman, 2014), SqueezeNet (Iandola et al., 2016), MobileNet (Howard et al., 2017), and DenseNet (Huang et al., 2017) feature extractors attached to an SSD detection head, with DenseNet yielding the best performance (Cordone et al., 2022). By designing a full-spike residual block, the capability to directly train deep-SNNs for object detection improved, outperforming hybrid models and achieving real-time responses (Su et al., 2023). Building on spiking residual blocks, Spiking-Retinanet proposed an ANN-SNN detector (Zhang et al., 2023), while Spiking-YOLOv4 was developed using a CNN-to-SNN method (Wang et al., 2023b). Additionally, an SNN version of a Region Proposal Network (RPN) for object recognition was introduced (Bulzomi et al., 2023). Spiking Fusion Object Detector (SFOD) was the first to adapt multi-scale feature fusion for SNNs using spiking data (Fan et al., 2024b). Recently, a framework integrating the entire process from event sampling to feature extraction in an end-to-end fashion achieved competitive results with ANNs (Wang et al., 2024).

The Recurrent Event-camera Detector (RED) uses Squeeze-and-Excitation (SE) layers (Hu et al., 2020) for feature extraction and Convolutional Long Short-Term Memory (ConvLSTM) blocks for spatiotemporal data extraction, combined with an SSD detection head (Perot et al., 2020). The Asynchronous Spatio-Temporal Memory Network (ASTMNet) comprises three components: Adaptive Temporal Sampling (ATS), Temporal Attention Convolutional Network (TACN), and Spatio-Temporal Memory. ATS samples events into bins based on an adaptive scheme related to the event frequency within an interval, while TACN aggregates events into an event representation called Asynchronous Attention Embedding. The Spatio-Temporal Memory module implements Recurrent-Convolutional (Rec-Conv) blocks following some convolutional layers (Li et al., 2022a). The Agile Event Detector (AED) introduced a new event representation called Temporal Active Focus (TAF) to encode sparse event streams into dense tensors, enhancing temporal information extraction (Liu et al., 2023). The Dual Memory Aggregation Network (DMANet) combines event information over different temporal ranges (short-term and long-term) with a

learnable representation, EventPillars, for the detection task (Wang et al., 2023a). A YOLOv5 (Jocher, 2020) detector was adapted to detect events encoded in a novel representation called Hyper Histograms, resulting in a remarkable reduction in terms of latency (Peng et al., 2023a).

Recurrent Vision Transformer (RVT) uses multi-axis attention (Tu et al., 2022) as a backbone, combined with ConvLSTMs (Shi et al., 2015) and YOLOX detection head for event-based detection (Gehrig and Scaramuzza, 2023). Enhancements to RVT through a self-labeling approach demonstrated further improvements (Wu et al., 2024). A detector based on SWin-v2 (Liu et al., 2022) was proposed, utilizing event encodings optimized through the Gromov-Wasserstein Discrepancy. This approach achieved state-of-the-art performance without the need for recurrent modules (Zubić et al., 2023). HMNet proposed a multi-rate hierarchy with multiple cells to model local and global context information from objects with varying dynamics, introducing sparse cross-attention operations between features and events (Hamaguchi et al., 2023). A transformer backbone featuring dual attention for spatial and polarity-temporal domains, paired with an event encoding focused on tokens, was also proposed (Peng et al., 2023b). To address event sparsity, a mechanism for processing only tokens with meaningful information was recently introduced, including a version of the Self-Attention operation adjusted for unequal token sizes (Peng et al., 2024). Recently, State Space Models (SSM) were introduced to replace RNN-cells for temporal modeling on detectors based on transformer backbones (Zubic et al., 2024).

Most of the aforementioned works process event features through some network and then adopt detection heads used on frame data, where the YOLO family is the most common choice. From this family, YOLOv8 is a well-acknowledged Object Detector in terms of performance, real-time operation, and scalability (Jocher et al., 2020). However, it works only with frames, which, in turn, lacks resources for processing event-based data, such as temporal-based processing. As mentioned before, a common solution for this is to add recurrent cells to frame-based extractors, as done in Perot et al. (2020), Li et al. (2022a), and Gehrig and Scaramuzza (2023), for example. Then, in this work, an extension of the YOLOv8 framework is proposed to add compatibility with events processing and training.

## 2.3 Data augmentation techniques for events

EventDrop randomly drops events from an event stream, which can be applied to individual events, events within a specific spatial location, or events within a particular time window (Gu et al., 2021). Neuromorphic Data Augmentation (NDA) introduced an augmentation policy incorporating techniques such as Horizontal Flip, Rolling, Rotation, CutOut, and CutMix for training SNNs (Li et al., 2022b). Spatio-temporal augmentation using random translation and time scaling was also proposed (Xiao et al., 2022). Temporal Event Shifting, which involves randomly reallocating events from a given frame to one of its prior frames, has proven beneficial for visual-aided force measurement (Naeini et al.,

2022). Event Spatiotemporal Fragments combines the inversion of event fragments on spatiotemporal and polarity domains with spatiotemporal drift of some slices of events through a certain extent (Shen et al., 2023b). Moreover, a viewpoint transform based on translation and rotation, combined with spatiotemporal stretching to prevent information loss due to out-of-resolution events discarded during the initial transformation, was introduced for training SNNs (Shen et al., 2023c).

EventCopyDrop is an enhanced version of EventDrop. It includes an additional augmentation called EventCopy, which creates copies of events from one random region and places them in another random location within the stream (Barchid et al., 2023). EventMix proposed an augmentation method based on mixing data from different event streams (Shen et al., 2023a). In RVT, Zoom-Out and Zoom-In augmentations were introduced to enhance event-based object detection (Gehrig and Scaramuzza, 2023). A framework combining geometric spatial augmentations with random temporal shifts and random polarity inversion was proposed (El Shair et al., 2023).

Shadow Mosaic is a technique that simulates events with varying densities, referred to as Shadows, and arranges them into a Mosaic to create a larger sample (Peng et al., 2023a). ShapeAug introduces random occlusions to event data, enhancing the robustness of object recognition applications (Bendig et al., 2024). Relevance Propagation Guidance (RPG) is employed to drop and mix events, resulting in the EventRPG augmentation method (Sun et al., 2024). EventAugment is an augmentation policy learning framework with 13 specific operations for event data, including flips, translations, crops, drops, and shear operations, targeting both spatial and temporal domains (Gu et al., 2024). Additionally, a temporal augmentation technique that involves dropping multiple sections of events within the temporal domain was proposed and evaluated for Lip-Reading applications (Dampfhofer and Mesquida, 2024).

Despite numerous works presenting various approaches for event data augmentation across spatial, temporal, and polarity domains, certain phenomena associated with event-based camera operations that can affect the generalization of Deep Learning models remain underexplored. In real-world scenes, brightness distribution is often irregular, and polarity distribution can vary significantly from one scene to another. Additionally, types of noise discussed in Lichtsteiner et al. (2008) and the adjustable bias settings in pixel circuits (Delbruck et al., 2010; Yang et al., 2012) can contribute to this variability. In this work, we propose a novel data augmentation model called Random Polarity Suppression to train Deep Learning models considering these variations.

### 3 Methodology

This work proposes an event-based object detection framework based on YOLOv8. To support this, a novel event data encoding method is presented, aiming to convert event streams into CNN-suitable representations that can be calculated with low latency, resulting in tensors that require low bandwidth and memory. Moreover, a data augmentation technique involving the random suppression of positive and negative polarities is also introduced to enhance the system's performance.

### 3.1 Volume of ternary event images

Event-based cameras function as 2D sensors capturing brightness variations at the pixel level. This process can be expressed mathematically as:

$$\Delta L(x_k, y_k, t_k) \geq p_k C \quad (1)$$

Here,  $\Delta L$  denotes the logarithmic change in a photoreceptor's current (brightness) at pixel location  $(x_k, y_k)$  and time  $t_k$ . The polarity  $p_k \in \{+1, -1\}$  indicates that a brightness change exceeding a threshold  $C$  in absolute value triggers a positive or negative event (Lichtsteiner et al., 2008). An event is characterized by the tuple  $e_k = (x_k, y_k, p_k, t_k)$ .

The sparse format of event streams poses a challenge for many current Deep Learning algorithms, requiring preprocessing for compatibility, as discussed in Section 2.1. With event-based cameras capable of operating at higher rates, driven by sensor resolution improvements (Gehrig and Scaramuzza, 2022), utilization of raw event data in downstream tasks can be complex. Efficient transformation involves sampling event streams at a consistent rate, partitioning them into sub-windows before conversion to dense tensors. This temporal binning strategy, effective in preserving temporal context from event streams, accommodates various conversion approaches such as applying a bilinear kernel to normalized timestamps (Zhu et al., 2018b; Perot et al., 2020), event counting (Gehrig and Scaramuzza, 2023), or tracking the latest event at particular locations (Nam et al., 2022). In this study, leveraging the success of such methodologies, a variation of Event Volumes is adopted, focusing on computation time and memory requirements. The chosen encoding scheme focuses on sampling the last event at each spatiotemporal location, as done in MDES (Nam et al., 2022), but with uniform temporal bin sizes.

Initially, a tensor  $I$  of dimensions  $B \times H \times W$  is initialized with zeros, where  $B$  represents the temporal bins and  $H, W$  are spatial dimensions. For a stream  $E = e_{k=0}^{N-1}$  containing  $N$  events sampled at a consistent time window  $T$ , a temporal division into uniform intervals is executed using the formula:

$$T_k = \lfloor \frac{t_k - t_0}{t_N - t_0} \rfloor B \quad (2)$$

In this equation,  $t_0$  represents the stream's initial timestamp,  $t_N$  the final timestamp, and  $T_k$  indicates the assigned temporal bin for the timestamp  $t_k$ . Subsequently, each pixel location  $(X_C, Y_C)$  within each temporal-based channel  $i \in 0, 1, \dots, B - 1$  is populated according to

$$I(X_C, Y_C, i) = \text{LastEvent}(E_{X_C, Y_C, i}) \quad (3)$$

where  $E_{X_C, Y_C, i} \in E$  represents the event subset associated with spatiotemporal location  $(X_C, Y_C, i)$ , and *LastEvent* extracts the polarity of the last event on it.

Unlike MDES (Nam et al., 2022), which generates grayscale bins, this method encodes the last event data using two values, +1 and -1, retaining the original polarity without alteration. Designating the background as 0, each bin corresponds to a

Ternary Event Image. Consequently, stacked bins form a Volume of Ternary Event Images (VTEI), as depicted in Figure 1. Employing only three values benefits memory-constrained contexts like edge applications. Moreover, the background data can be discarded for compression purposes, minimizing bandwidth requirements for data transmission. A swift computation is ultimately anticipated, as the primary operation involves mapping each event to its relevant position in the dense grid.

### 3.2 Recurrent YOLOv8 architecture

The architecture of Recurrent YOLOv8 (ReYOLOv8) is specifically designed to leverage the spatio-temporal properties of event-based data. Its baseline is inspired by the YOLOv8 family of models, which are well-regarded for their efficiency and accuracy in real-time frame-based object detection tasks (Jocher et al., 2020). Previous studies have also documented the successful application of other YOLO versions in event-based object detection (Gehrig and Scaramuzza, 2023; Zubić et al., 2023). Moreover, research indicates that incorporating recurrent connections into models enhances their efficiency, allowing for better handling of sequential data and improving overall performance (Gehrig and Scaramuzza, 2023; Perot et al., 2020). Similarly to YOLOv8, ReYOLOv8 consists of a backbone, neck, and detection head. The primary distinction between the YOLOv8 and ReYOLOv8 models lies in replacing the original C2f block in the architecture with a Recurrent C2f block, making the backbone recurrent. As it can be seen from Figure 2, the incoming event streams converted to VTEI tensors with five bins using the method outlined in Section 3.1, are fed to the recurrent backbone of ReYOLOv8. Its first layers are comprised of feature extractors, structured as in the original YOLOv8 (Jocher et al., 2020), but incorporating recurrent blocks and resizing certain convolutional blocks. In particular, the *Conv2D* blocks function as standard convolutional layers for spatial feature downsampling. Starting from the 2nd stage, feature maps pass through *Rec C2f* blocks for further refinement before downsampling. *Rec C2f* blocks combine *C2f* blocks adopted from YOLOv8 model with a *ConvLSTM* (Shi et al., 2015) block. This combination enables feature refinement within the channel domain while capturing long-range temporal relationships by leveraging both current and past features. Subsequent to the final recurrent block, a Spatial Pyramid Pooling (SPP) block enriches features by combining multiple receptive fields (He et al., 2015). The final features produced by this Recurrent Backbone are fed into YOLOv8's neck, Path Aggregation Network (PANet) (Liu et al., 2018), for further fusion and transmission to the 3-level detection heads.

Figure 3 illustrates the architecture of a *C2f* block (Jocher et al., 2020), which is an efficient version of a Cross-Stage Partial (CSP) Bottleneck block (Bochkovskiy et al., 2020) with two convolutions. The initial convolution in this block adjusts the input channel count. Subsequently, a *Split* block separates the feature into two groups with equal channels. One group undergoes processing through a sequence of *N* Bottleneck blocks, with the same structure as the ResNet's blocks (He et al., 2016). Notably, the shortcut connections within these blocks are deactivated

when incorporated into the PANET framework. Finally, the remaining split channels are merged with the output from each Bottleneck, followed by another convolution to reduce the number of channels.

The *ConvLSTM* block depicted in Figure 2 implements a Long-Short Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997) cell using  $1 \times 1$  convolutions instead of the traditional Multi-Layer Perceptrons (MLP). This adaptation maintains the fundamental operational principle of an LSTM while enhancing its versatility to accommodate varying spatial dimensions. A similar structure was adopted on RVT (Gehrig and Scaramuzza, 2023). When considering an input  $x$ , along with the previous hidden  $h_{t-1}$  and cell states  $c_{t-1}$  of the block, the functionality of this cell can be expressed through the following equations:

$$i = \sigma(\text{Conv}2D_{1 \times 1}([x, h_{t-1}])) \quad (4)$$

$$r = \sigma(\text{Conv}2D_{1 \times 1}([x, h_{t-1}])) \quad (5)$$

$$o = \sigma(\text{Conv}2D_{1 \times 1}([x, h_{t-1}])) \quad (6)$$

$$c = \tanh(\text{Conv}2D_{1 \times 1}([x, h_{t-1}])) \quad (7)$$

In these equations,  $i$ ,  $r$ ,  $o$ , and  $c$  represent the input, remember, output, and cell gates, respectively. Here,  $\text{Conv}2D_{1 \times 1}$  symbolizes a 2D convolution with a  $1 \times 1$  kernel,  $\sigma$  denotes a sigmoid activation function, and the square brackets indicate concatenation between two inputs. The current hidden and cell states,  $h_t$  and  $c_t$ , respectively, can be determined by:

$$c_t = r * c_{t-1} + c * i \quad (8)$$

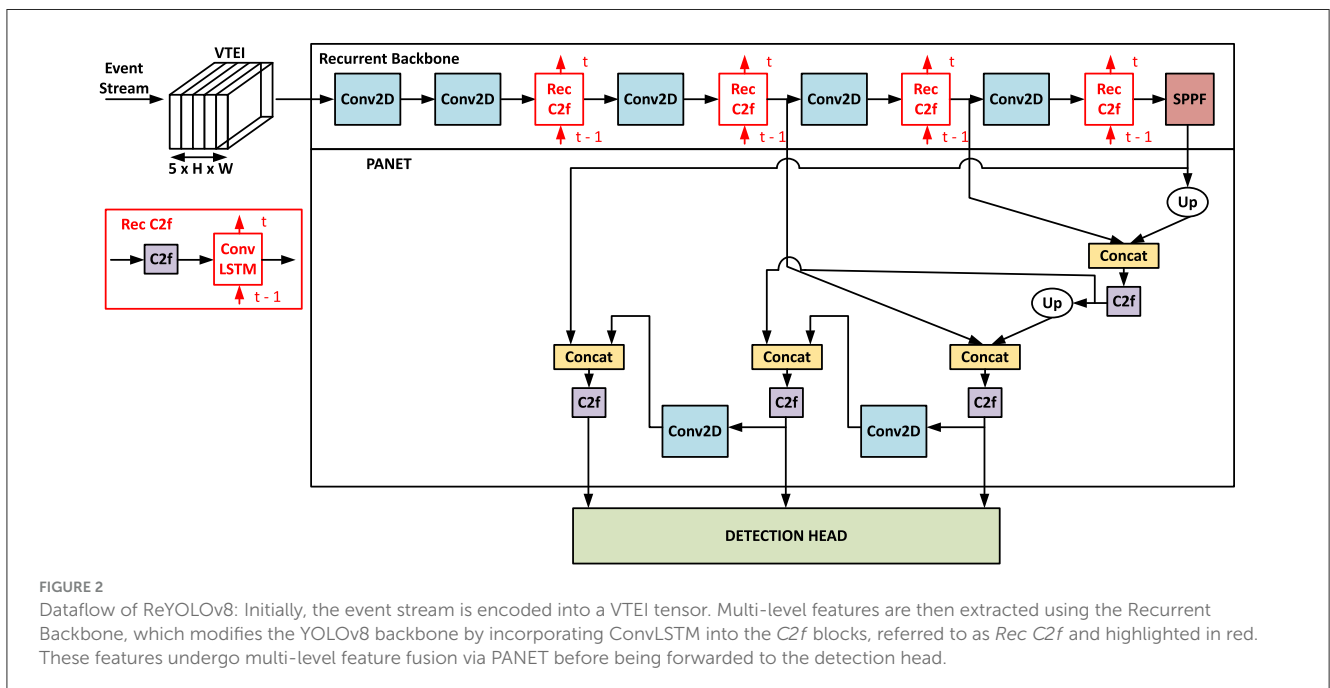
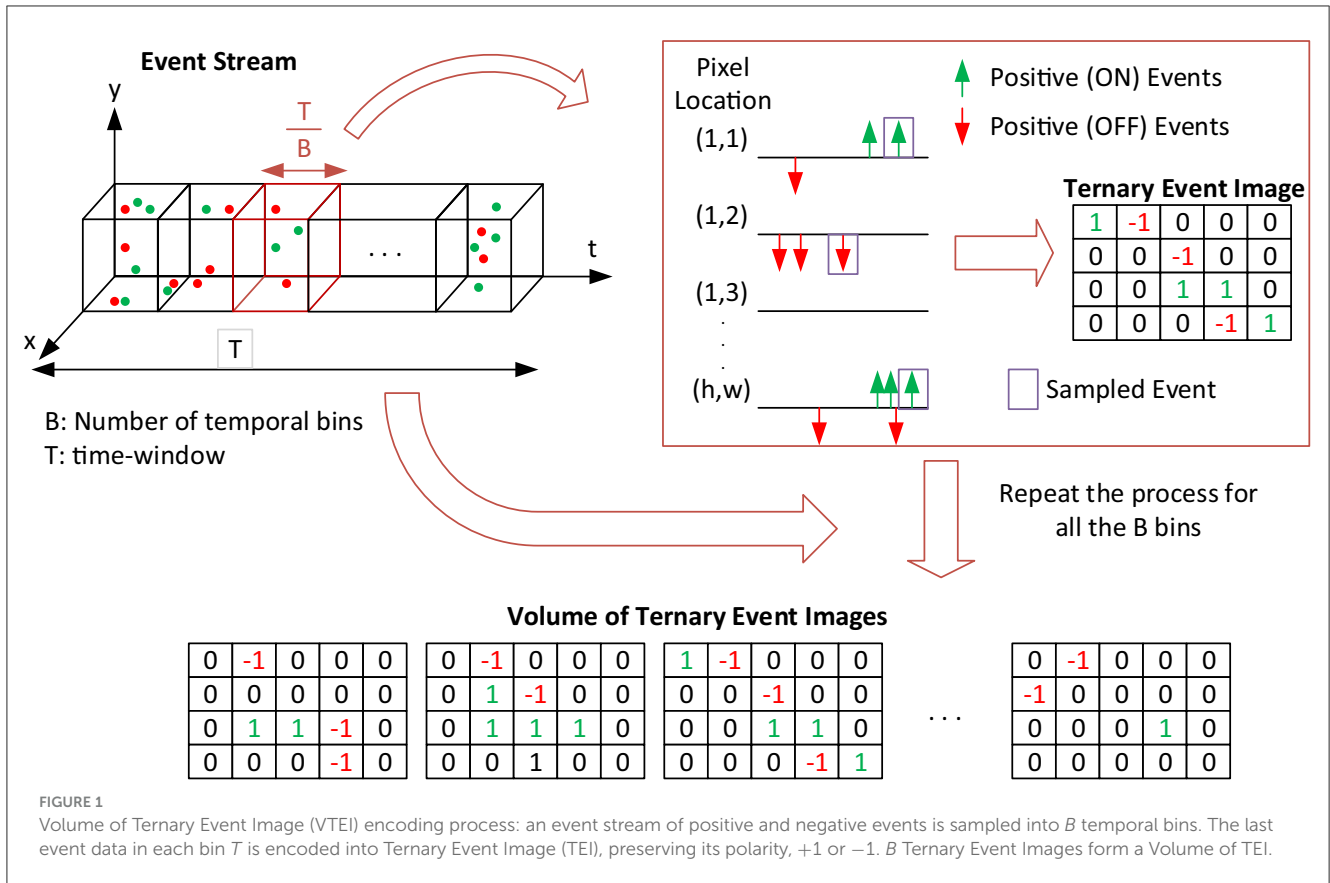
$$h_t = o * \tanh(c_t) \quad (9)$$

Additionally,  $h_t$  is the output feature to be propagated to the subsequent block.

This work introduces three variants of the Recurrent YOLOv8 architecture displayed in Figure 2: ReYOLOv8n, ReYOLOv8s, and ReYOLOv8m, representing nano, small, and medium scales, respectively. These scales align with the standards set by the original YOLOv8 (Jocher et al., 2020). Alongside incorporating recurrent cells, modifications are made regarding the channel count in each layer and the number of bottleneck blocks within each *C2f*, which are slightly adjusted compared to the reference models. Table 1 provides a comprehensive overview of the modules within ReYOLOv8s, where  $C_{in}$  and  $C_{out}$  denote the number of input and output channels, respectively. ReYOLOv8n and ReYOLOv8m follow a similar framework, implementing distinct channel and width multipliers to rescale the number of channels and bottlenecks. In this sense, new models are created in a way similar to Compound Scaling (Tan and Le, 2019). All *Conv2D* operations are utilized for downsampling, utilizing a kernel size of 3 and a stride of 2.

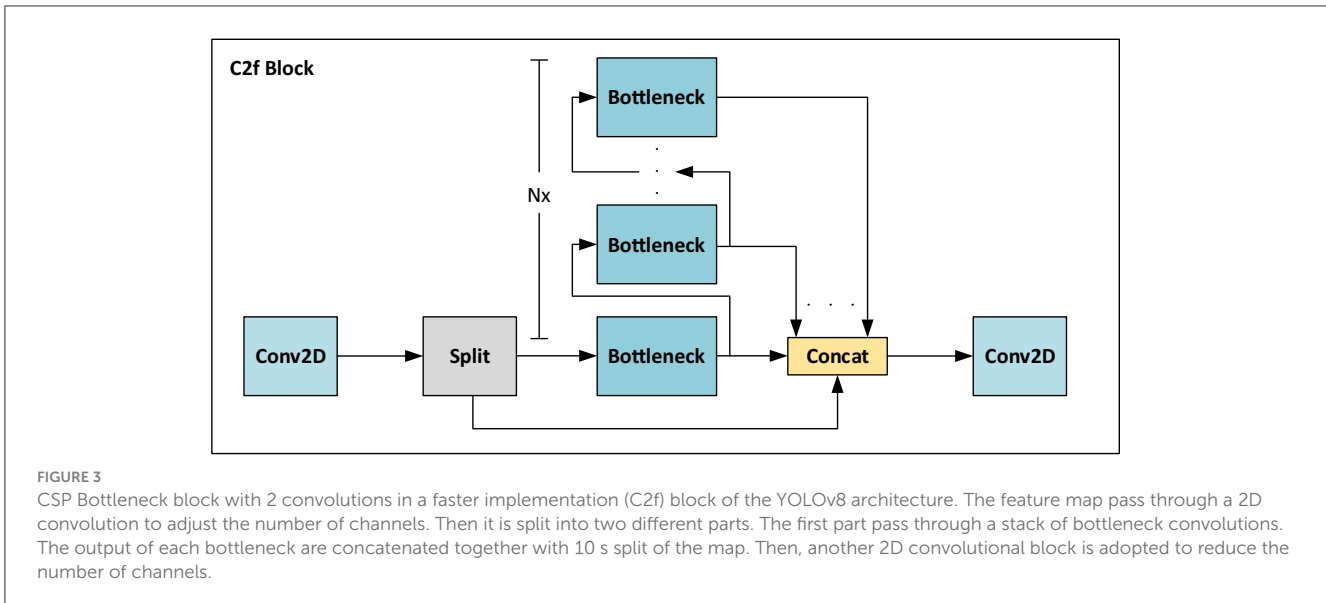
### 3.3 Event data augmentation with random polarity suppression

The polarity imbalance in event data can stem from various factors. First and foremost, changes in illumination



are scene-dependent, making it difficult to ensure an equal distribution of positive and negative events corresponding to scene movements. Furthermore, electronic circuits within pixels are susceptible to noise across multiple stages, ranging from the photoreceptor sensor to the comparator and amplifier stages

(Posch et al., 2014). Even when other noise-related parameters are well-controlled, sporadic positive polarity noisy events have been documented (Lichtsteiner et al., 2008). Moreover, bias currents within different stages of an event-based camera pixel can be externally adjusted, potentially influencing sensitivity to distinct



polarities in varying ways (Delbruck et al., 2010; Yang et al., 2012).

Taking this into consideration, a data augmentation technique that specifically targets the polarity domain is proposed. To train the detector effectively under unbalanced polarity scenarios, a probability  $s$  will be introduced to suppress a specific polarity within each batch. Additionally, another probability  $p$  will denote the likelihood of suppressing the positive polarities, with  $(1 - p)$  representing the corresponding value for the negative ones. Considering the random variables  $(r1, r2) \in [0, 1]$ , the subset of all pixels from the VTEI tensor  $I$  with negative polarity  $I_n$ , and its positive counterpart  $I_p$ , the Random Polarity Suppression (RPS) technique will construct a new tensor  $I'$  based on the following condition:

$$I' = \begin{cases} I, & \text{if } r1 \geq s, \text{ else} \\ I_p, & \text{if } r2 \geq p, \text{ else} \\ I_n, & \text{otherwise} \end{cases} \quad (10)$$

For event encodings like VTEI, where there is a direct correlation between tensor values and their original event stream polarity, this augmentation can be applied post-conversion, reducing complexity compared to applying it to raw data. To ensure the consistency of the augmentation, the same modifications are applied across all temporal bins. Figure 4 provides a schematic illustrating how RPS works on a given VTEI tensor, accompanied by grayscale representations of such images.

Given that this work utilizes VTEI encoding, the proposed augmentation technique reduces all information associated with a particular polarity. However, for encodings dependent on event counting, such as Stacked Histograms (Gehrig and Scaramuzza, 2023), it is feasible to implement this augmentation more gradually. This approach involves reducing the content of a specific polarity to a certain degree rather than entirely erasing it. To demonstrate the effectiveness and robustness of this technique,

it is tested on real-world and extensive datasets known for their complexity in scenes and lighting conditions. Such datasets present a more challenging environment for evaluating the technique's performance and adaptability.

### 3.4 Datasets

The object detection models in this work were validated using two substantial real-world event datasets, as outlined in Table 2. The first dataset, PEDRo, is designed for person detection with a primary focus on Robotics applications. Recorded in Italy using a handheld camera, PEDRo captures individuals across diverse scenes, lighting conditions, and weather situations. The data was captured using a DAVIS346 camera with a resolution of  $346 \times 260$  pixels. PEDRo is the sole real-world event-based large-scale dataset tailored specifically for Robotics applications to date (Boretti et al., 2023). The second dataset, Prophese's Generation 1 Automotive Dataset (GEN1), was recorded in France and encompasses various weather and illumination scenarios incorporating pedestrians and cars (De Tournemire et al., 2020). While both datasets are significant in their respective applications and sizes, they exhibit complementary characteristics. GEN1 boasts a wider range of object classes; however, there exists an imbalance of  $\sim 5:1$  between cars and pedestrians. Moreover, the pedestrian class is predominantly represented on smaller scales and toward the sides of the images, aligning with the expected viewpoint from a car's perspective. Conversely, PEDRo offers a more uniform representation of pedestrians across the pixel grid and showcases a greater diversity in terms of aspect ratios compared to GEN1 (Boretti et al., 2023). Additionally, for GEN1, box filtering was applied to remove bounding boxes with one of the diagonals smaller than 30 or one of the sides smaller than 10, as done by Perot et al. (2020).



TABLE 1 The structure of ReYOLOv8n, ReYOLOv8s, ReYOLOv8m models.

Node	Input Nodes	Block Name	ReYOLOv8n			ReYOLOv8s			ReYOLOv8m			Residual (C2f only)
			Cin	Cout	Bottlenecks (C2f only)	Cin	Cout	Bottlenecks (C2f only)	Cin	Cout	Bottlenecks (C2f only)	
<b>Recurrent backbone</b>												
1	–	Conv2D	5	16	–	5	24	–	5	32	–	–
2	1	Conv2D	16	32	–	24	48	–	32	64	–	–
3	2	C2f	32	32	2	48	48	2	64	64	2	–
4	3	ConvLSTM	32	43	–	48	48	–	64	64	–	–
5	4	Conv2D	32	64	–	48	88	–	64	128	–	–
6	5	C2f	64	64	3	88	88	3	128	128	3	–
7	6	ConvLSTM	64	64	–	88	88	–	128	128	–	–
8	7	Conv2D	64	128	–	88	176	–	128	256	–	–
9	8	C2f	128	128	3	176	176	3	256	256	3	–
10	9	ConvLSTM	128	128	–	176	176	–	256	256	–	–
11	10	Conv2D	128	256	–	176	344	–	256	512	–	–
12	11	C2f	256	256	2	344	344	2	512	512	2	–
13	12	ConvLSTM	256	256	–	344	344	–	512	512	–	–
14	13	SPFF	256	256	–	344	344	–	512	512	–	–
<b>PANET</b>												
15	14	Upsample	256	256	–	344	344	–	512	512	–	–
16	15,10	Concat	344,128	384	–	344,176	520	–	512,256	768	–	–
17	16	C2f	384	128	2	520	176	2	768	256	2	False
18	17	Upsample	128	128	–	176	176	–	256	256	–	–
19	17,7	Concat	128,64	192	–	176,88	264	–	256,128	384	–	–
20	19	C2f	192	64	2	264	88	2	384	128	2	False
21	20	Conv2D	64	64	–	88	88	–	128	128	–	–
22	21,17	Concat	64,128	192	–	88,176	264	–	128,256	384	–	–
23	22	C2f	192	128	2	264	176	2	384	256	2	False
24	23	Conv2D	128	128	–	176	176	–	256	256	–	–
25	24,14	Concat	128,256	384	–	176,344	520	–	256,512	768	–	–
26	25	C2f	384	256	2	520	344	2	768	512	2	False
<b>Detection head</b>												
27	20,23,26	Detect	64,128,256	–	–	88,176,344	–	–	128,256,512	–	–	–

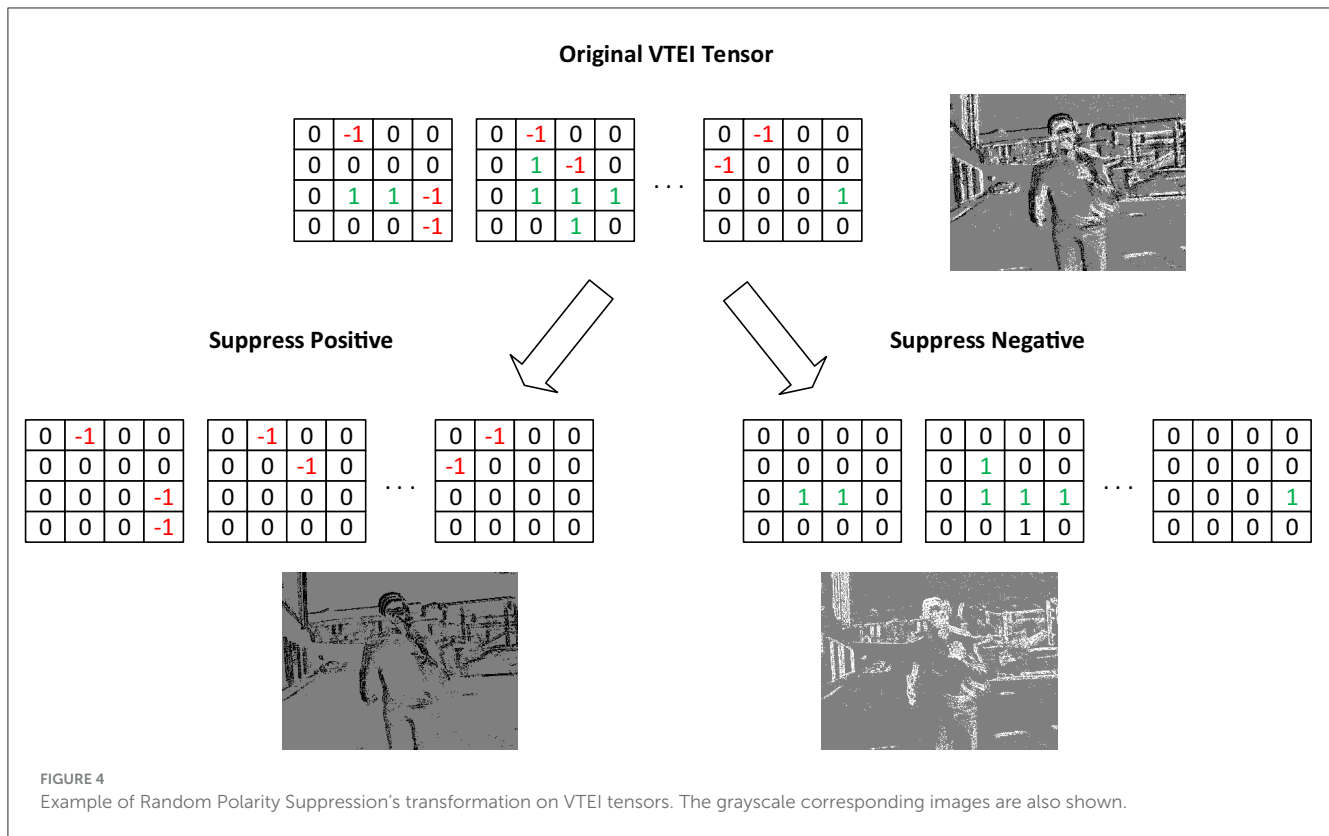


TABLE 2 Datasets adopted in this work.

Dataset	Focus	Classes	Resolution	Labels
PEDRO (Boretti et al., 2023)	Robotics	1	346 × 260	43k
GEN1 (De Tourneville et al., 2020)	Automotive	2	304 × 240	255k

### 3.5 Training and evaluation procedure

To train the models described in Section 2.2, Truncated Backpropagation Through Time (T-BPTT) (Werbos, 1990) was employed. During training, each dataset was segmented into clips with limited sequence lengths, and the memory cells were reset after each clip. During validation, complete original sequences were assessed, with memory cells being reset at the end of each sequence, aligning with methodologies observed in Perot et al. (2020) and Gehrig and Scaramuzza (2023). Consistent application of data augmentation techniques was ensured across all frames within the same training sequence. The optimizer utilized was Stochastic Gradient Descent (SGD), with a momentum of 0.937 and linear learning rate decay.

In addition to T-BPTT, the training process closely adhered to the approach established within the YOLOv8 framework. A warm-up phase of three epochs was adopted to initiate training, consisting of a momentum of 0.8 and a bias learning rate of

0.1. The losses for box regression, classification, and Distribution Focal Loss (DFL) (Li et al., 2020) maintained the same values as the original framework: 7.5, 0.5, and 1.5, respectively. Dataset-specific hyperparameters are detailed in Table 3. All models were trained from scratch for 100 epochs. The sequence length is related to the T-BPTT setup. The image sizes from both datasets were adjusted to multiples of 32, aligning with the YOLOv8 anchors (Jocher et al., 2020).  $LR_0$  and  $LR_f$  reference the initial and final learning rates, respectively, determined by the prescribed schedule. Probability values for HFlip and Zoom-Out denote the likelihood of horizontal flip and zoom-out augmentations, with zoom-out scales randomly chosen between 1.0 and 1.2 for all cases. To speed up the training, validations were run every 10 epochs.

All training procedures were conducted utilizing a V100 GPU, while speed validations were executed on NVIDIA GTX1080ti and V100 GPUs for literature comparison. The entire development process was done using the PyTorch library (Ansel et al., 2024). The primary evaluation metric employed in this study was the Microsoft-Common Objects in Context (MS-COCO) mean Average Precision (mAP), specified as  $mAP@0.5:0.95$ . For this metric, a confidence threshold of 0.0001, consistent with YOLOv8, was utilized, along with a Non-Maximum Suppression (NMS) threshold of 0.7. During speed calculations, the confidence threshold was set to 0.25, and the Intersection over Union (IoU) was defined at 0.45. The final performance results were determined by evaluating the test subsets with the models that presented the highest mAP for the val subsets for each dataset.

TABLE 3 Hyperparameters adopted for training the models for GEN1 and PEDRo datasets.

Dataset	Model	Sequence Length	Image Size	Batch Size	Weight Decay	LRO	LRF	HFlip	Zoom Out
GEN1	ReYOLOv8n	11	320 × 256	48	0.011	0.01	0.0001	0.5	0.5
	ReYOLOv8s	11	320 × 256	48	0.011	0.01	0.0001	0.5	0.5
	ReYOLOv8m	11	320 × 256	36	0.011	0.01	0.0007	0.5	0.5
PEDRO	ReYOLOv8n	5	352 × 288	48	0.005	0.07	0.0007	0.5	0.2
	ReYOLOv8s	5	352 × 288	48	0.005	0.07	0.0007	0.5	0.2
	ReYOLOv8m	5	352 × 288	48	0.005	0.07	0.0007	0.5	0.2

## 4 Results and discussion

### 4.1 Evaluation of VTEI

In order to assess the effectiveness of encoding event streams using VTEI, the test set from the GEN1 dataset was converted into three distinct formats commonly found in Event-Based Object Detection literature: Voxel Grids (Zhu et al., 2018b), MDES (Nam et al., 2022), and Stacked Histograms (SHist) (Gehrig and Scaramuzza, 2023). The specific variation of Voxel Grids utilized by RED was employed (Perot et al., 2020). The comparison excluded formats such as Hyper Histograms (Peng et al., 2023a), an extended version of SHist with additional channels; Event Temporal Images (Fan et al., 2024a), similar to VTEI but with distinct accumulation and mapping processes; and asynchronous formats like TORE (Baldwin et al., 2022) and TAF tensors (Liu et al., 2023), due to focusing on fixed time window encodings. Additionally, format-specific to different model categories like Voxel Cubes (Cordone et al., 2022) for SNNs, Group Tokens (Peng et al., 2023b) for transformers, and graph-related representations (Bi et al., 2019; Schaefer et al., 2022; Gehrig and Scaramuzza, 2024) were also excluded, as this study is centered on formats that can be adopted alongside convolutional architectures.

Table 4 provides dimensions and the minimum byte counts for VTEI, SHist, and MDES for the GEN1 dataset. Compared to MDES and VTEI, SHist and Voxel Grids calculate polarities in separate channels before stacking them temporally, effectively doubling the number of channels. For GEN1, dimensions  $B, H,$  and  $W$  are 5, 240, and 304, respectively. The minimum byte count for each format was computed based on the Coordinate List (COO) compression, where each non-zero is encoded by its coordinate and data content. For the GEN1, spatial dimensions are encoded in 17 bits. The data content for VTEI and MDES can be represented in binary format, while the number of bins can be encoded in three bits. Then, summing it up, VTEI and MDES require three bytes per non-zero entry. SHist needs an additional bit to represent the channel domain, while the data itself is encoded in eight bits, totaling four bytes per entry. Voxel Grids is similar to SHist but utilizes a sum function over normalized timestamps instead of event counting. Hence, with a decimal range requirement, a half-precision floating-point format (16 bits) is required, which is the smallest present in libraries like Pytorch, adding an extra byte for data encoding. Overall, the merging of polarities into the same channels and employing a narrower data

TABLE 4 Comparison between different event encodings in dimensions and minimum number of bytes required to encode data from the GEN1 dataset.

Data format	Dimensions	Minimum # bytes
VTEI (this work)	$B \times H \times W$	3
SHist (Gehrig and Scaramuzza, 2023)	$2B \times H \times W$	4
MDES (Nam et al., 2022)	$B \times H \times W$	3
Voxel grid (Zhu et al., 2018b; Perot et al., 2020)	$2B \times H \times W$	5

The calculations consider  $B = 5, H = 240,$  and  $W = 304.$

value range positioned VTEI and MDES as superior options in terms of memory requirements.

To better highlight the capabilities of each format, a further analysis was performed, taking into consideration the recording from the GEN1's test set with the biggest number of events. All analyses were conducted on 50 ms samples extracted from the original recordings. The outcomes, summarized in Table 5, are categorized into three sections detailing the average, maximum, and minimum event counts across all chunks. The performance data were gathered using an Intel Xeon Gold 6230R CPU on an Ubuntu 20.04.6 LTS operating system with 251 GB of RAM. Regarding latency for a 50 ms window (Latency@50ms), based on the value most commonly adopted in the literature, VTEI emerged as the top performer compared to other formats in all scenarios. Notably, VTEI was  $1.53\times$  faster on average than SHist. Comparatively, VTEI showcased an average  $6.73\times$  speed-up over MDES, being  $5.10\times$  faster at the peak event number, while these metrics were  $3.4\times$  and  $3.0\times$  compared to Voxel Grids, respectively. The Event Rate in Mev/s, reflecting the event processing capabilities, also favored VTEI across all scenarios, with an average rate  $\sim 3.89\times$  higher than alternative formats. The event rate increased  $2.13\times$  from the minimum to the maximum event counts, highlighting the scalability of VTEI.

This table also reveals the number of non-zero elements after encoding for the three scenarios. The "Encoded Size" field represents the disk space required to encode each scenario based on the number of bytes shown in Table 4. The Compression Ratio is derived by dividing the disk size of the reference events in each section by their respective Encoded Sizes. It can be seen that VTEI achieves a compression ratio ranging from  $2.53\times$

TABLE 5 Comparison between different event encodings based on chunks of event streams with 50 ms from the test set of the GEN1 dataset in terms of speed, compression rate, and bandwidth.

Data format	Latency @50ms	Event rate (Mev/s)	Non-zeros	Encoded size (MB)	Compression Ratio	BW (MB/s)
<b>Scenario 1, Average #Events processed = 192,063 (0.73MB)</b>						
VTEI	<b>1.5 ms</b>	<b>128.24</b>	96,017	0.27	2.67	5.33
SHist	2.3 ms	83.63	125,234	0.48	1.53	9.14
MDES	10.1 ms	18.96	<b>89,532</b>	<b>0.26</b>	<b>2.86</b>	<b>4.26</b>
Voxel grid	5.1 ms	37.87	125,232	0.60	1.23	10.84
<b>Scenario 2, Maximum #Events processed = 327,020 (1.25MB)</b>						
VTEI	<b>2.2 ms</b>	<b>145.69</b>	172,654	0.49	2.53	9.45
SHist	3.4 ms	95.22	219,848	0.84	1.49	15.70
MDES	11.2 ms	29.32	<b>145,305</b>	<b>0.42</b>	<b>3.00</b>	<b>6.80</b>
Voxel grid	6.6 ms	49.49	219,846	1.05	1.19	18.52
<b>Scenario 3, Minimum #Events processed = 44,962 (0.17MB)</b>						
VTEI	<b>0.6 ms</b>	<b>68.43</b>	20,310	0.058	2.95	1.15
SHist	1.1 ms	40.62	29,268	0.11	1.54	2.18
MDES	2.1 ms	21.49	<b>19,832</b>	<b>0.057</b>	<b>3.02</b>	<b>1.09</b>
Voxel grid	2.2 ms	20.42	29,266	0.14	1.23	2.67

The bold values represent the best performance.

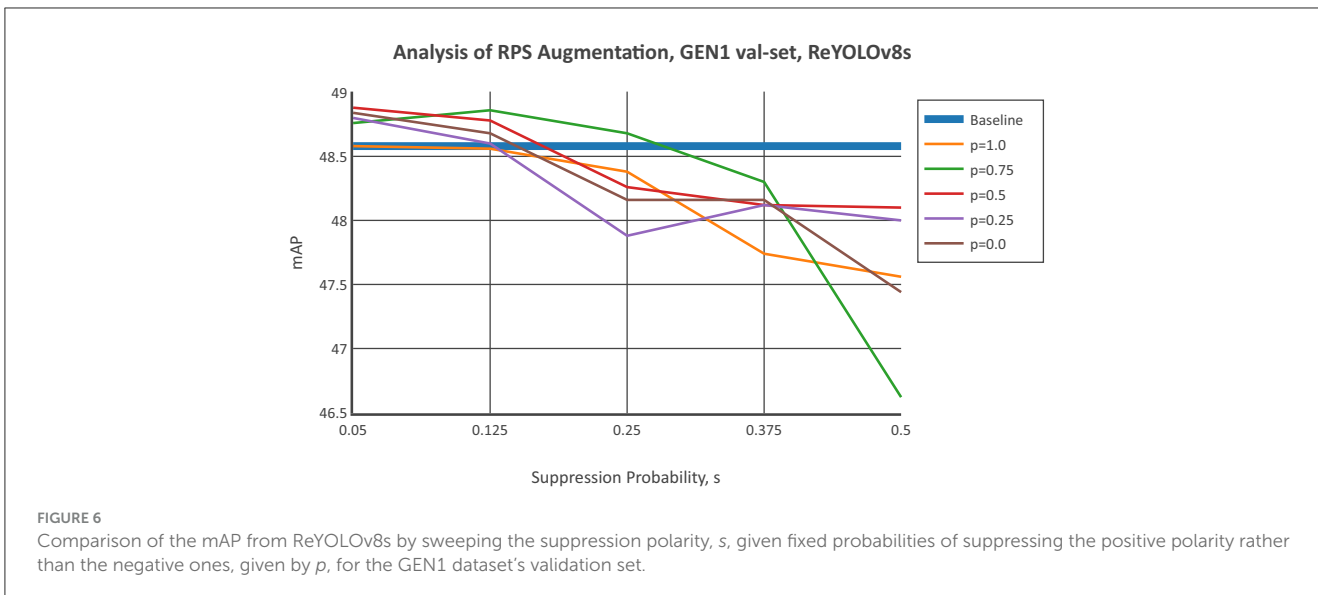
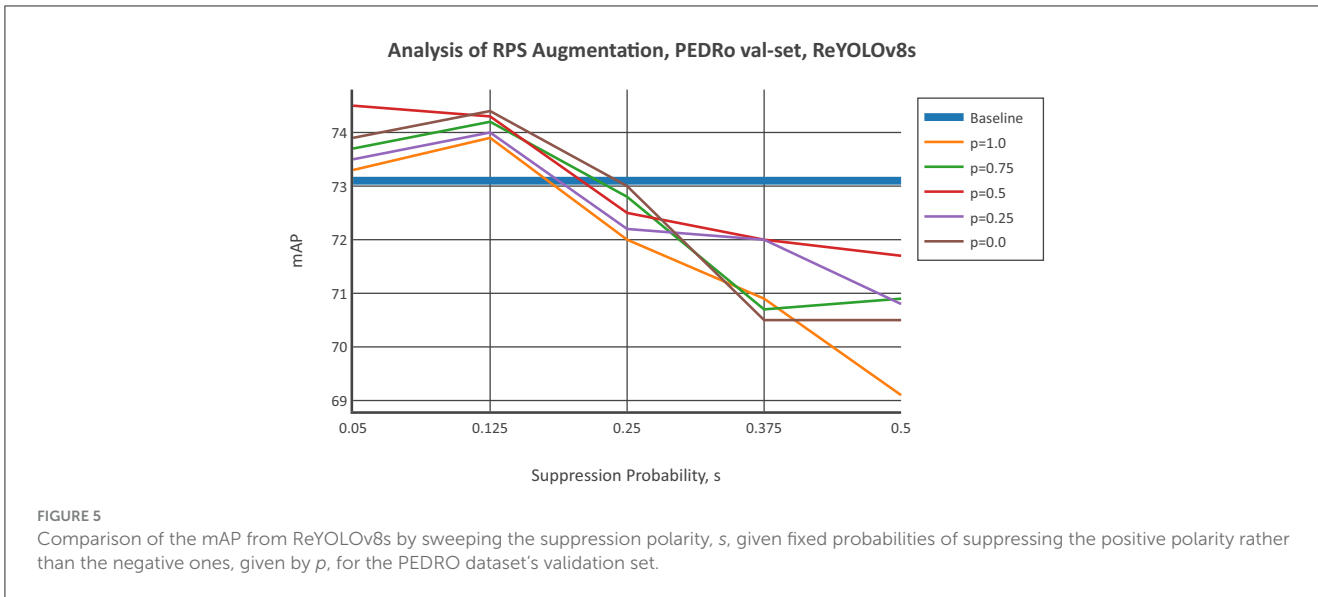
to  $2.95\times$ , which is, on average,  $1.75\times$  higher than SHist,  $2.0\times$  better than Voxel Grids, and  $\sim 93.4\%$  as efficient as MDES. It is worth mentioning that MDES adopts non-uniform temporal bins, whose lengths decrease by powers of two, where some intervals are larger than those seen with VTEI. By sampling only the latest events in each bin, more events tend to be ignored at larger intervals. The larger sub-interval of MDES comprises half of the whole time window, contributing to a smaller number of non-zero elements and related metrics. Additionally, to provide communication-related insights, a Bandwidth (BW) field is included, calculated by dividing the Encoded Sizes by the sum of the total sampling period (assumed to be 50 ms here) and the latencies from the conversion process. In this context, VTEI's bandwidth ranges from 1.15 to 9.45 MB/s, which is, on average,  $1.7\times$  lower than SHist,  $2.0\times$  better than Voxel Grid, and only  $1.25\times$  higher than MDES. In conclusion, VTEI stands out as the best choice for latency and event processing rate among the options presented, requiring the lowest number of bytes for data encoding (alongside MDES), and being competitive with MDES regarding memory requirements.

## 4.2 Evaluation of random polarity suppression

The impact of applying Polarity Suppression (RPS) on models was evaluated by comparing the baseline mean Average Precision (mAP) of ReYOLOv8s with the results using RPS. The positive polarity suppression probability  $p$  was initially set at five levels: 0.0, 0.25, 0.50, 0.75, 1.0. Different suppression probabilities  $s$  from 0.05, 0.125, 0.25, 0.375, 0.5 were tested for each  $p$  value.

Five runs were conducted for both the baseline and each RPS combination, and average results were reported to account for variability. Figure 5 displays the validation set results from the PEDRo dataset. The findings show that improvements generally occur up to a 12.5% suppression probability, suggesting that RPS enhances model performance when used as a small disturbance. Larger suppression levels may degrade the quality of training samples and negatively impact performance. The most significant enhancements across all  $p$  values were seen at  $s = 0.05$  and  $s = 0.125$ . Notably, on average,  $p = 0.50$  led to the greatest improvements, followed by  $p = 0.0$ ,  $p = 0.75$ ,  $p = 0.25$ , and  $p = 1.0$ . There is no clear pattern regarding which polarity should be suppressed, as similar performance levels were observed for  $p = 0.0$ , which represents a total negative polarity suppression, and  $p = 0.75$ , an aggressive positive polarity suppression. A comparable trend was seen for  $p = 0.25$  and  $p = 1.0$ , which involve opposite types of data manipulation but resulted in similar outcomes at  $s = 0.125$ . These results suggest that objects in the PEDRo dataset have a diverse polarity distribution and are not particularly biased toward one polarity. This also explains why the most balanced scenario,  $p = 0.50$ , led to the best results on average.

Figure 6 shows the results of applying RPS to the GEN1 validation set, using the same procedure as for PEDRo. A similar trend to Figure 5 is observed, where improvements in mean Average Precision (mAP) are more significant for suppression probabilities below 12.5%. Except for  $p = 0.75$ , no improvements were observed for  $s = 0.25$ . When examining positive and negative suppression separately, the balanced scenario with  $p = 0.50$  performed better than most other values, similar to the findings for PEDRo. However, an exception was observed for  $p = 0.75$ , which achieved the highest average improvement and a peak



enhancement comparable to  $p = 0.50$ . At  $p = 0.0$  and  $p = 0.25$ , there was a noticeable increase in mAP at  $s = 0.05$ , with steep decreases after  $s = 0.125$ . No improvement was observed for  $p = 1.0$ .

Table 6 expands on the previous analysis to include other scales of ReYOLOv8, detailing the combinations of suppression and positive/negative probabilities that resulted in the most significant improvements for each case. Due to time constraints, single runs were conducted for this analysis. Since more runs were performed for the ReYOLOv8s models, the baseline and RPS cases with the highest mAPs on each validation set were selected. When individually assessing the models, it was observed that GEN1 showed comparable improvement across all scales. In contrast, for PEDRo, the Nano and Medium models exhibited more substantial improvements compared to the Small variant. Interestingly, a common trend across all models pointed toward

greater enhancements at lower suppression probabilities, consistent with the patterns identified in Figures 5, 6. Moreover, a balanced probability of positive and negative suppression was usually the most effective. Variations in scenes due to changes in illumination or camera instability, which contribute to polarity imbalance, along with other factors highlighted in Section 4.2, can impact the efficacy of the suppression technique on object detection performance. The stochastic nature of data augmentation, where samples are randomly subjected to transformations like RPS, further complicates predicting whether suppressing positive or negative polarities will yield improvements. After suppression, the modified tensors may not align with the intended polarity distribution. In conclusion, the analysis suggests that maintaining a sufficiently low level of general suppression may lead to enhancements at different positive and negative suppression levels, with a balanced probability being the safest choice for the latter.

TABLE 6 Analysis of the effect of Random Polarity Suppression over the mAP of the models presented in this work.

Dataset	Model	Suppression	Positive	Baseline mAP	mAP with RPS	Improvement
GEN1	ReYOLOv8n	0.25	0.25	45.9	46.3	+0.4
	ReYOLOv8s	0.05	0.50	48.1	48.3	+0.2
	ReYOLOv8m	0.05	0.50	49.0	49.4	+0.4
PEDRo	ReYOLOv8n	0.125	0.25	59.0	63.9	+4.9
	ReYOLOv8s	0.05	0.50	64.5	65.5	+1.0
	ReYOLOv8m	0.125	0.50	66.5	69.1	+2.6

TABLE 7 Comparison with the literature for the PEDRo dataset.

Model	Network	Parameters	mAP	Runtime
ReYOLOv8n (this work)	CNN + RNN	4.7 M	63.9	<b>9.2 ms</b>
ReYOLOv8s (this work)	CNN + RNN	8.4 M	65.5	10.4 ms
ReYOLOv8m (this work)	CNN + RNN	18.1 M	<b>69.1</b>	12.3 ms
YOLOv8x (Boretti et al., 2023)	CNN	68.2 M	58.6	17.6 ms

The bold values represent the best performance.

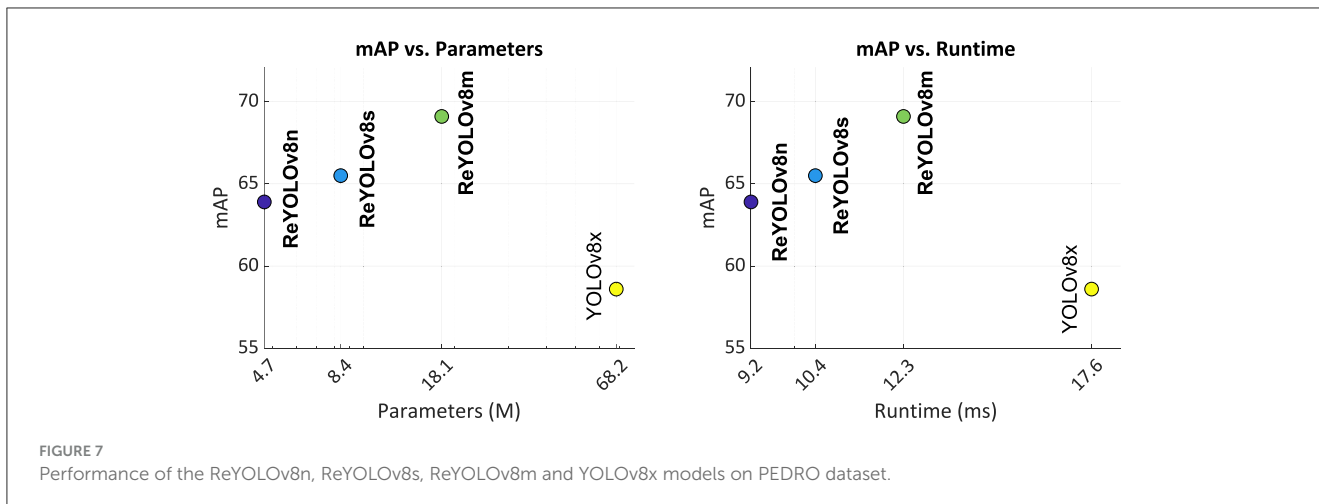
### 4.3 Comparison with the state-of-the-art

Table 7 and Figure 7 compare the three models introduced in this study, namely ReYOLOv8n, ReYOLOv8s, and ReYOLOv8m, with the state-of-the-art YOLOv8x-based model for PEDRo. All models exhibited improvements in mean Average Precision (mAP), ranging from 9% to 18%, requiring significantly fewer parameters—14.5× and 3.8×, respectively. This notable performance improvement at a lower number of parameters can be attributed to the integration of long-range temporal modeling in the models of this study, a feature lacking in the benchmark YOLOv8x-based model. Given that the original work did not include runtime information (Boretti et al., 2023), the inference times for YOLOv8x processing the same VTEI tensors as ReYOLOv8 were reported here. From this comparative analysis, the models in this study demonstrated an average speed-up of 1.46×, highlighting their efficiency. Furthermore, as shown in Table 6, the ReYOLOv8n implementation without RPS achieved a similar mAP to YOLOv8x, differing by only 0.6%. However, after applying RPS, the gap between the models widened to 9%. This indicates that data augmentation through RPS significantly contributed to the improvements observed in this study, in addition to utilizing memory cells.

Table 8 and Figure 8 present a comparative analysis of the state-of-the-art models for the GEN1 dataset. Models are classified into Nano, Small, and Medium scales based on their trainable parameter ranges (5 M, 15 M, and 45 M, respectively), a common practice in Computer Vision literature. Inference times are reported separately for NVIDIA's V100 GPU and a group consisting of closely related devices: NVIDIA's GTX1080ti, Titan

XP, and GTX980. The results are illustrated in Figure 9, where only models with mAP exceeding 40.0 are considered, and the best results are highlighted. Firstly, analyzing the Nano models, ReYOLOv8n achieves a 5% improvement in mAP compared to RVT-T, requiring only an additional 0.3M parameters while reducing latency by 0.20 ms. When comparing the Small models, ReYOLOv8s also demonstrates superior performance, with a 2.8% increase in mAP compared to HMNet-L1, demanding around 27% fewer parameters but being 2.4× slower. On the medium scale, ReYOLOv8m outperforms SAST-CB (Peng et al., 2024) by 2.5%, but also with fewer parameters. The superior performance of ReYOLOv8 can be attributed to its foundation on the YOLOv8 baseline, known to outperform other detectors like YOLOX (Ge et al., 2021), which serves as the detection head for RVT (Gehrig and Scaramuzza, 2023), GET (Peng et al., 2023b), and SAST-CB (Peng et al., 2024) models. Models exceeding 45M parameters were omitted, as they did not yield substantial gains compared to the Medium-scaled models. Notable exceptions include DSTDNet-X (Fan et al., 2024a), which has 100 M parameters and has a mAP similar to ReYOLOv8m, and ERGO12 (Zubić et al., 2023), a model based on a pre-trained SWinv2 (Liu et al., 2022) transformer that achieves a mAP of 50.4 with 59.6 M parameters and a latency of 77 ms.

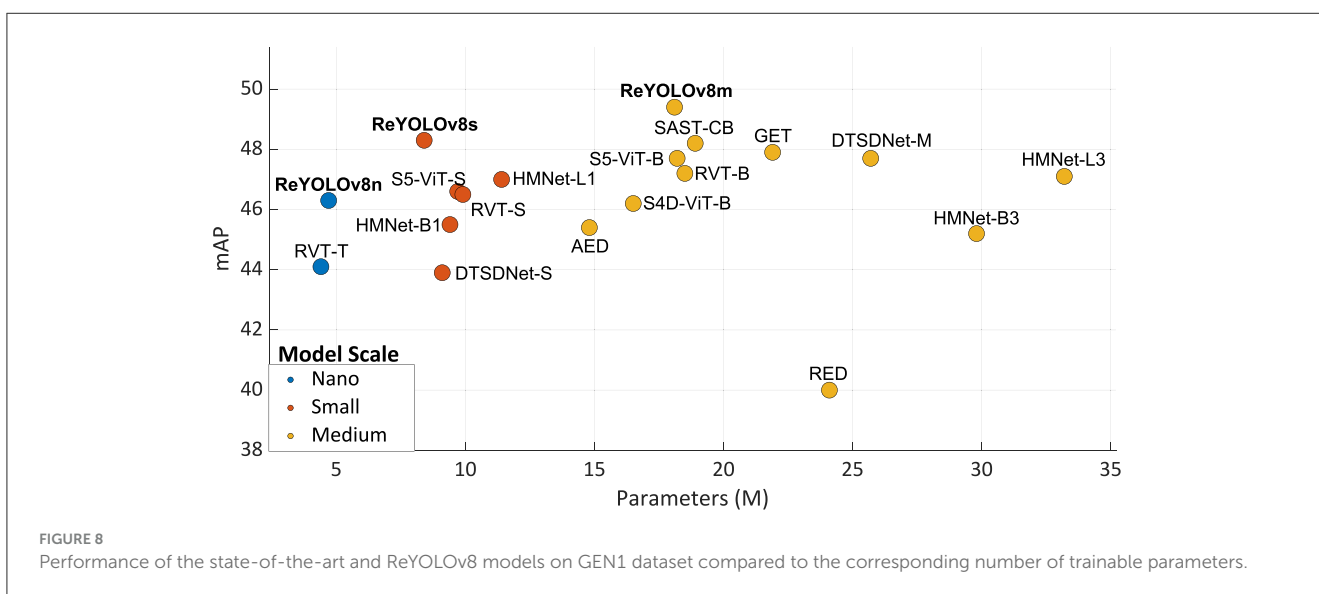
In terms of runtime, it can be seen that the models are competitive when compared to other approaches that deploy RNNs, where ReYOLOv8n is 2% faster than RVT-T, RVT-S outperforms ReYOLOv8s by around 9%, and ReYOLOv8m achieves 80% of the speed of RVT-B. On the other hand, compared to models that do not deploy RNNs, ReYOLOv8s is, on average, 1.9× slower than other similar-sized models. A similar behavior is also present for Medium models. A deeper look into their design choices should be done to understand this difference. For the case of DTSDNet (Fan et al., 2024a), events are converted into two different tensors according to different time-window sizes, a smaller and a bigger one. Those tensors are passed through two parallel CNNs, and the resulting features are fused afterward. By doing so, they can introduce long- and short-term temporal modeling without adopting RNNs, avoiding the latency those cells introduce. AED (Liu et al., 2023) achieves a similar effect by creating tensors with long-range temporal information incremented periodically through updates of a FIFO-queue, while HMNET (Hamaguchi et al., 2023) leverages a hierarchical structure with latent memories working at different rates. Furthermore, replacing RNN cells with SSM also led to speed-ups (Zubic et al., 2024). However, those models are generally outperformed by RNN-based models, where

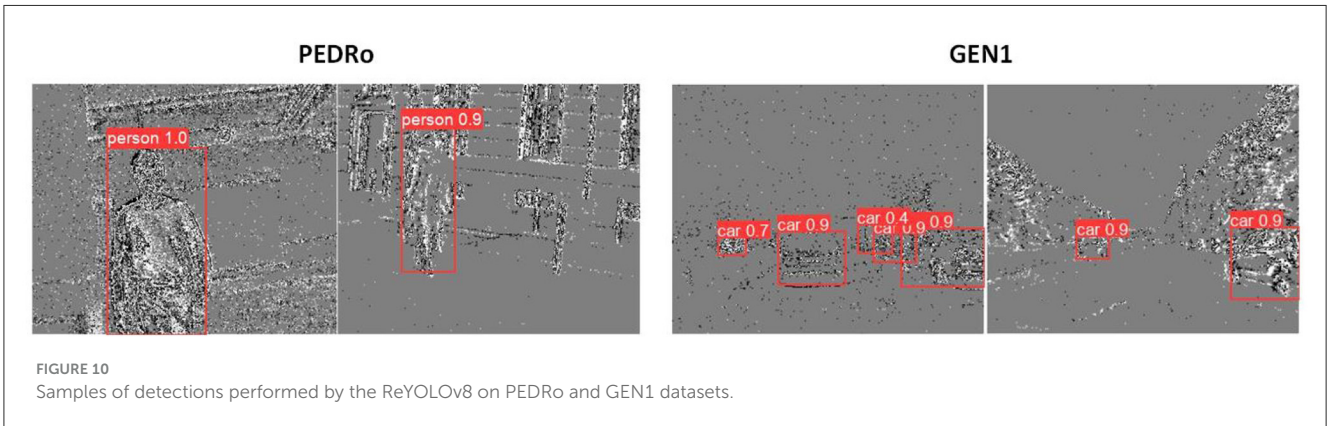
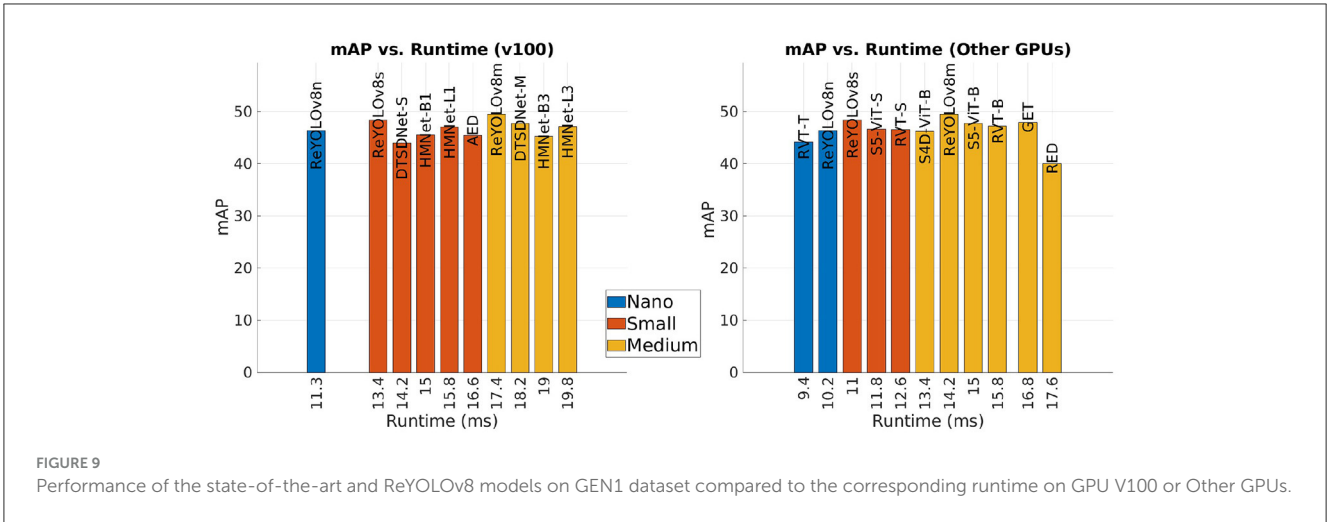


**TABLE 8** Comparison between the different parameters for the GEN1 dataset present in the literature.

	Model	Network	Parameters	mAP	Runtime (ms)
Nano	RVT-T (Gehrig and Scaramuzza, 2023)	Transformer + RNN	4.4 M	44.1	9.4 ms
	RYOLOv8n (this work)	CNN + RNN	4.7 M	<b>46.4</b>	<b>9.2 ms</b>
Small	RYOLOv8s (this work)	CNN + RNN	<b>8.4 M</b>	<b>48.6</b>	10.4 ms
	DTSDNet-S (Fan et al., 2024a)	Dual Pathway CNN	9.1 M	43.9	-
	RVT-S (Gehrig and Scaramuzza, 2023)	Transformer + RNN	9.9 M	46.5	<b>9.5 ms</b>
Medium	RYOLOv8m (this work)	CNN + RNN	<b>18.1M</b>	<b>49.5</b>	12.3 ms
	GET (Peng et al., 2023b)	Transformer	18.4 M	47.9	16.8 ms
	RVT-B (Gehrig and Scaramuzza, 2023)	Transformer + RNN	18.5 M	47.2	<b>10.2 ms</b>
	SAST-CB (Peng et al., 2024)	Transformer + RNN	18.9 M	48.2	-
	RED (Perot et al., 2020)	CNN + RNN	24.1 M	40.0	16.7 ms
	DTSDNet-M (Fan et al., 2024a)	Dual Pathway CNN	25.7 M	47.7	-

The bold values represent the best performance.





at a Small scale, ReYOLOv8s has a 2.8% better mAP than HMNet-L1, with 1.4× fewer parameters and a 3.6% better performance when compared to S5-ViT-S, which is based on SSM cells. In contrast, at the Medium scale, ReYOLOv8m has a 3.6% higher performance when compared to DTSDNet-M, with a 1.42× smaller model. ReYOLOv8m also has an average of 5.2% higher mAP than the SSM-based models S4D-ViT-B and S5-ViT-B.

In conclusion, by integrating recurrent cells and leveraging a robust baseline, the ReYOLOv8 models achieve higher mAP than other models of similar scales. Although they incur some latency penalties compared to non-RNN models, they require fewer parameters. However, the latencies ranging from 9.2 to 15.5 ms are still suitable for real-time operation. Figure 10 shows some samples of detections performed by the ReYOLOv8 framework.

## 5 Conclusion

In this work, we proposed a framework for event-based object detection by modifying the frame-based detector YOLOv8 to include recurrent cells, resulting in the new Recurrent YOLOv8 (ReYOLOv8) model. Additionally, we introduced novel event encoding and event-specific data augmentation techniques. The proposed event encoding, called Volume of Ternary Event Images

(VTEI), leverages the spatio-temporal dynamics of event streams while minimizing processing time, memory requirements, and bandwidth. VTEI was shown to be 3.21× faster than alternatives when processing peak event loads, requiring only 3 bytes per non-zero and a peak bandwidth of 9.45 BM/s.

Three scale-based variations of ReYOLOv8 Nano, Small, and Medium were evaluated using two large-scale, real-world datasets: PEDRo and GEN1. To address polarity imbalances in event streams, we proposed a data augmentation technique that randomly suppresses specific polarities, leading to an average mAP improvement of 0.7% for GEN1 and 4.5% for PEDRo. Compared to existing literature, the ReYOLOv8 models demonstrated a 9%–18% improvement in mAP for PEDRo, while requiring 8.8× fewer parameters and being 1.67× faster on average. For GEN1, the models presented outstanding performance across all the similar scales, achieving mAP improvements of ~5%, 2.8%, and 2.5% over the runner-ups up to 5 M, 15 M, and 45 M parameters, respectively, with an average parameter reduction of 20.8%. The inference speeds, ranging from 9.2 to 15.5 ms, make these models suitable for real-time operation.

Future work will extend the Random Suppression Polarity analysis to other event encodings. Additionally, there is a need to create comprehensive benchmarks to better understand the system-level impacts of jointly designing event encodings and detection models, as this remains underexplored. Finally, we intend



to broaden the framework evaluation to include other datasets available in the literature, including 1MegaPixel dataset (Perot et al., 2020).

## Data availability statement

The original contributions presented in the study are included in the article/supplementary material and the codes are available on GitHub at the following link <https://github.com/silvada95/ReYOLOv8>. Further inquiries can be directed to the corresponding author.

## Author contributions

DS: Conceptualization, Investigation, Methodology, Software, Validation, Writing – original draft. KS: Conceptualization, Investigation, Validation, Writing – review & editing. AE: Investigation, Methodology, Validation, Writing – review & editing. MF: Conceptualization, Methodology, Project administration, Supervision, Writing – review & editing. AME: Funding acquisition, Supervision, Writing – review & editing.

## Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. This work has been funded by King Abdullah University of Science and Technology CRG program under grant number: URF/1/4704-01-01.

## References

- Annamalai, L., Chakraborty, A., and Thakur, C. S. (2021). EvAn: neuromorphic event-based sparse anomaly detection. *Front. Neurosci.* 15:699003. doi: 10.3389/fnins.2021.699003
- Ansel, J., Yang, E., He, H., Gimelshein, N., Jain, A., Voznesensky, M., et al. (2024). "PyTorch 2: faster machine learning through dynamic python bytecode transformation and graph compilation," in *29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24)* (New York, NY: ACM). doi: 10.1145/3620665.3640366
- Baldwin, R. W., Almatrafi, M., Kaufman, J. R., Asari, V., and Hirakawa, K. (2019). "Inceptive event time-surfaces for object classification using neuromorphic cameras," in *Image Analysis and Recognition: 16th International Conference, ICIAR 2019, Waterloo, ON, Canada, August 27-29, 2019, Proceedings, Part II 16* (Cham: Springer), 395–403. doi: 10.1007/978-3-030-27272-2\_35
- Baldwin, R. W., Liu, R., Almatrafi, M., Asari, V., and Hirakawa, K. (2022). Time-ordered recent event (TORE) volumes for event cameras. *IEEE Trans. Pattern Anal. Mach. Intell.* 45, 2519–2532. doi: 10.1109/TPAMI.2022.3172212
- Barchid, S., Mennesson, J., and Djéraba, C. (2023). "Exploring joint embedding architectures and data augmentations for self-supervised representation learning in event-based vision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Vancouver, BC: IEEE), 3903–3912. doi: 10.1109/CVPRW59228.2023.00405
- Bendig, K., Schuster, R., and Stricker, D. (2024). Shapeaug: occlusion augmentation for event camera data. *arXiv [Preprint]*. arXiv:2401.02274. doi: 10.48550/arXiv.2401.02274
- Benosman, R., Clercq, C., Lagorce, X., Ieng, S.-H., and Bartolozzi, C. (2013). Event-based visual flow. *IEEE Trans. Neural Netw. Learn. Syst.* 25, 407–417. doi: 10.1109/TNNLS.2013.2273537
- Bi, Y., Chadha, A., Abbas, A., Boursoulatte, E., and Andreopoulos, Y. (2019). "Graph-based object classification for neuromorphic vision sensing," in *Proceedings of the IEEE/CVF international conference on computer vision* (Seoul: IEEE), 491–501. doi: 10.1109/ICCV.2019.00058
- Bochkovskiy, A., Wang, C.-Y., and Liao, H.-Y. M. (2020). Yolov4: Optimal speed and accuracy of object detection. *arXiv [Preprint]*. arXiv:2004.10934. doi: 10.48550/arXiv.2004.10934
- Boretti, C., Bich, P., Pareschi, F., Prono, L., Rovatti, R., Setti, G., et al. (2023). "Pedro: an event-based dataset for person detection in robotics," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Vancouver, BC: IEEE), 4064–4069. doi: 10.1109/CVPRW59228.2023.00426
- Bulzomi, H., Gruel, A., Martinet, J., Fujita, T., Nakano, Y., Bendahan, R., et al. (2023). "Object detection for embedded systems using tiny spiking neural networks: Filtering noise through visual attention," in *2023 18th International Conference on Machine Vision and Applications (MVA)* (Hamamatsu: IEEE), 1–5. doi: 10.23919/MVA57639.2023.10215590
- Chen, G., Cao, H., Conradt, J., Tang, H., Rohrbein, F., Knoll, A., et al. (2020). Event-based neuromorphic vision for autonomous driving: a paradigm shift for bio-inspired visual sensing and perception. *IEEE Signal Process. Mag.* 37, 34–49. doi: 10.1109/MSP.2020.2985815
- 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)* (Padua: IEEE), 1–8. doi: 10.1109/IJCNN55064.2022.9892618
- Dampfhofer, M. and Mesquida, T. (2024). "Neuromorphic lip-reading with signed spiking gated recurrent units," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Seattle, WA: IEEE), 2141–2151. doi: 10.1109/CVPRW63382.2024.00219

## Acknowledgments

We express our sincere gratitude to the KAUST Supercomputing Laboratory for granting us access to its GPU clusters. Additionally, we are deeply thankful to Mathias Gehrig from the Robotics and Perception Group at the University of Zurich for his invaluable insights and information on the implementation aspects of recurrent-convolutional architectures and managing large-scale event datasets. We acknowledge the use of ChatGPT-4o to enhance the clarity and readability of this manuscript.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be considered a potential conflict of interest.

The author(s) declared that they were an editorial board member of Frontiers, at the time of submission. This had no impact on the peer review process and the final decision.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- 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. doi: 10.48550/arXiv.2001.08499
- Delbruck, T., Berner, R., Lichtsteiner, P., and Dualibe, C. (2010). "32-bit configurable bias current generator with sub-off-current capability," in *Proceedings of 2010 IEEE International Symposium on Circuits and Systems (Paris: IEEE)*, 1647–1650. doi: 10.1109/ISCAS.2010.5537475
- El Shair, Z. A., Hassani, A., and Rawashdeh, S. A. (2023). Cstr: a compact spatio-temporal representation for event-based vision. *IEEE Access* 11, 102899–102916. doi: 10.1109/ACCESS.2023.3316143
- 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 Trans. Intell. Veh.* 1–15. doi: 10.1109/TIV.2024.3365991
- Fan, Y., Zhang, W., Liu, C., Li, M., and Lu, W. (2024b). Sfd: spiking fusion object detector. *arXiv [Preprint]*. arXiv:2403.15192. doi: 10.48550/arXiv.2403.15192
- Gallego, G., Delbruck, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., et al. (2022). Event-based vision: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 154–180. doi: 10.1109/TPAMI.2020.3008413
- Ge, Z., Liu, S., Wang, F., Li, Z., and Sun, J. (2021). YoloX: exceeding yolo series in 2021. *arXiv [Preprint]*. arXiv:2107.08430. doi: 10.48550/arXiv.2107.08430
- Gehrig, D., Loquercio, A., Derpanis, K. G., and Scaramuzza, D. (2019). "End-to-end learning of representations for asynchronous event-based data," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (Seoul: IEEE)*, 5633–5643. doi: 10.1109/ICCV.2019.00573
- Gehrig, D. and Scaramuzza, D. (2022). Are high-resolution event cameras really needed? *arXiv [Preprint]*. arXiv:2203.14672. doi: 10.48550/arXiv.2203.14672
- Gehrig, D. and Scaramuzza, D. (2024). Low-latency automotive vision with event cameras. *Nature* 629, 1034–1040. doi: 10.1038/s41586-024-07409-w
- 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 (Vancouver, BC: IEEE)*, 13884–13893. doi: 10.1109/CVPR52729.2023.01334
- Gu, F., Dou, J., Li, M., Long, X., Guo, S., Chen, C., et al. (2024). EventAugment: learning augmentation policies for asynchronous event-based data. *IEEE Trans. Cogn. Dev. Syst.* 16, 1521–1532. doi: 10.1109/TCDS.2024.3380907
- Gu, F., Sng, W., Hu, X., and Yu, F. (2021). EventDrop: data augmentation for event-based learning. *arXiv [Preprint]*. arXiv:2106.05836. doi: 10.48550/arXiv.2106.05836
- 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 (Vancouver, BC: IEEE)*, 22867–22876. doi: 10.1109/CVPR52729.2023.02190
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 37, 1904–1916. doi: 10.1109/TPAMI.2015.2389824
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (Las Vegas, NV: IEEE)*, 770–778. doi: 10.1109/CVPR.2016.90
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780. doi: 10.1162/neco.1997.9.8.1735
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., et al. (2017). Mobilenets: efficient convolutional neural networks for mobile vision applications. *arXiv [Preprint]*. arXiv:1704.04861. doi: 10.48550/arXiv.1704.04861
- Hu, J., Shen, L., Albanie, S., Sun, G., and Wu, E. (2020). Squeeze-and-excitation networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 42, 2011–2023. doi: 10.1109/TPAMI.2019.2913372
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. (2017). "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition (Honolulu, HI: IEEE)*, 4700–4708. doi: 10.1109/CVPR.2017.243
- Iandola, F. N., Han, S., Moskewicz, M. W., Ashraf, K., Dally, W. J., Keutzer, K., et al. (2016). Squeezenet: alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size. *arXiv [Preprint]*. arXiv:1602.07360. doi: 10.48550/arXiv.1602.07360
- Jha, S., Seo, C., Yang, E., and Joshi, G. P. (2021). Real time object detection and tracking system for video surveillance system. *Multimed. Tools Appl.* 80, 3981–3996. doi: 10.1007/s11042-020-09749-x
- Jocher, G. (2020). YOLOv5 by Ultralytics. Available at: <https://github.com/ultralytics/yolov5> (accessed December 31, 2024).
- Jocher, G., Chaurasia, A., and Qiu, J. (2020). YOLO by Ultralytics. Available at: <https://github.com/ultralytics/ultralytics> (accessed December 31, 2024).
- Kugele, A., Pfeil, T., Pfeiffer, M., and Chicca, E. (2021). "Hybrid SNN-ANN: energy-efficient classification and object detection for event-based vision," in *DAGM German Conference on Pattern Recognition (Cham: Springer)*, 297–312. doi: 10.1007/978-3-030-92659-5\_19
- Lagorce, X., Ieng, S.-H., Clady, X., Pfeiffer, M., and Benosman, R. B. (2015). Spatiotemporal features for asynchronous event-based data. *Front. Neurosci.* 9:46. doi: 10.3389/fnins.2015.00046
- Li, J., Li, J., Zhu, L., Xiang, X., Huang, T., Tian, Y., et al. (2022a). Asynchronous spatio-temporal memory network for continuous event-based object detection. *IEEE Trans. Image Process.* 31, 2975–2987. doi: 10.1109/TIP.2022.3162962
- Li, X., Wang, W., Wu, L., Chen, S., Hu, X., Li, J., et al. (2020). Generalized focal loss: learning qualified and distributed bounding boxes for dense object detection. *Adv. Neural Inf. Process. Syst.* 33, 21002–21012.
- Li, Y., Kim, Y., Park, H., Geller, T., and Panda, P. (2022b). "Neuromorphic data augmentation for training spiking neural networks," in *European Conference on Computer Vision (Cham: Springer)*, 631–649. doi: 10.1007/978-3-031-20071-7\_37
- Lichtsteiner, P., Posch, C., and Delbruck, T. (2008). A 128 × 128 120 db 15  $\mu$ s latency asynchronous temporal contrast vision sensor. *IEEE J. Solid-State Circuits* 43, 566–576. doi: 10.1109/JSSC.2007.914337
- Lin, T., Goyal, P., Girshick, R., He, K., and Dollar, P. (2020). Focal loss for dense object detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 42, 318–327. doi: 10.1109/TPAMI.2018.2858826
- Liu, B., Xu, C., Yang, W., Yu, H., and Yu, L. (2023). Motion robust high-speed lightweight object detection with event camera. *IEEE Trans. Instrum. Meas.* 72:5013113. doi: 10.1109/TIM.2023.3269780
- Liu, L., Ouyang, W., Wang, X., Fieguth, P., Chen, J., Liu, X., et al. (2020). Deep learning for generic object detection: a survey. *Int. J. Comput. Vis.* 128, 261–318. doi: 10.1007/s11263-019-01247-4
- Liu, S., Qi, L., Qin, H., Shi, J., and Jia, J. (2018). Path aggregation network for instance segmentation. in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8759–8768. doi: 10.1109/CVPR.2018.00913
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., et al. (2016). "SSD: single shot multibox detector," in *Computer Vision-ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I 14 (Cham: Springer)*, 21–37. doi: 10.1007/978-3-319-46448-0\_2
- Liu, Z., Hu, H., Lin, Y., Yao, Z., Xie, Z., Wei, Y., et al. (2022). "Swin transformer v2: scaling up capacity and resolution," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (New Orleans, LA: IEEE)*, 12009–12019. doi: 10.1109/CVPR52688.2022.01170
- 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 (Salt Lake City, UT: IEEE)*, 5419–5427. doi: 10.1109/CVPR.2018.00568
- Messikommer, N., Gehrig, D., Loquercio, A., and Scaramuzza, D. (2020). "Event-based asynchronous sparse convolutional networks," in *Computer Vision-ECCV 2020: 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VIII 16 (Cham: Springer)*, 415–431. doi: 10.1007/978-3-030-58598-3\_25
- Michaelis, C., Mitzkus, B., Geirhos, R., Rusak, E., Bringmann, O., Ecker, A. S., et al. (2019). Benchmarking robustness in object detection: autonomous driving when winter is coming. *arXiv [Preprint]*. arXiv:1907.07484. doi: 10.48550/arXiv.1907.07484
- Naeini, F. B., Kachole, S., Muthusamy, R., Makris, D., and Zweiri, Y. (2022). Event augmentation for contact force measurements. *IEEE Access* 10, 123651–123660. doi: 10.1109/ACCESS.2022.3224584
- 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 (New Orleans, LA: IEEE)*, 6114–6123. doi: 10.1109/CVPR52688.2022.00602
- Peng, Y., Li, H., Zhang, Y., Sun, X., and Wu, F. (2024). "Scene adaptive sparse transformer for event-based object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (Seattle, WA: IEEE)*, 16794–16804. doi: 10.1109/CVPR52733.2024.01589
- Peng, Y., Zhang, Y., Xiao, P., Sun, X., and Wu, F. (2023a). Better and faster: adaptive event conversion for event-based object detection. *Proc. AAAI Conf. Artif. Intell.* 37, 2056–2064. doi: 10.1609/aaai.v37i2.25298
- Peng, Y., Zhang, Y., Xiong, Z., Sun, X., and Wu, F. (2023b). "Get: group event transformer for event-based vision," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (Paris: IEEE)*, 6038–6048. doi: 10.1109/ICCV51070.2023.00555
- Perot, E., De Tournemire, P., Nitti, D., Masci, J., and Sironi, A. (2020). Learning to detect objects with a 1 megapixel event camera. *Adv. Neural Inf. Process. Syst.* 33, 16639–16652.
- Posch, C., Serrano-Gotarredona, T., Linares-Barranco, B., and Delbruck, T. (2014). Retinomorphic event-based vision sensors: bioinspired cameras with spiking output. *Proc. IEEE* 102, 1470–1484. doi: 10.1109/JPROC.2014.2346153
- Rebecq, H., Horstschaefer, T., and Scaramuzza, D. (2017). *Real-Time Visual-Inertial Odometry for Event Cameras Using Keyframe-Based Nonlinear Optimization*. Available at: [https://www.zora.uzh.ch/id/eprint/139471/1/BMV/C17\\_Rebecq.pdf](https://www.zora.uzh.ch/id/eprint/139471/1/BMV/C17_Rebecq.pdf)
- 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 (New Orleans, LA: IEEE), 12371–12381. doi: 10.1109/CVPR52688.2022.01205
- Shen, G., Zhao, D., and Zeng, Y. (2023a). Eventmix: an efficient data augmentation strategy for event-based learning. *Inf. Sci.* 644:119170. doi: 10.1016/j.ins.2023.119170
- Shen, H., Luo, Y., Cao, X., Zhang, L., Xiao, J., Wang, T., et al. (2023b). “Training robust spiking neural networks on neuromorphic data with spatiotemporal fragments,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Rhodes Island: IEEE), 1–5. doi: 10.1109/ICASSP49357.2023.10096951
- Shen, H., Xiao, J., Luo, Y., Cao, X., Zhang, L., Wang, T., et al. (2023c). “Training robust spiking neural networks with viewpoint transform and spatiotemporal stretching,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Rhodes Island: IEEE), 1–5. doi: 10.1109/ICASSP49357.2023.10097016
- Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., and Wong, W.-K. Woo, W.-c. (2015). Convolutional lstm network: a machine learning approach for precipitation nowcasting. *Adv. Neural Inf. Process. Syst.* 28.
- Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv [Preprint]*. arXiv:1409.1556. doi: 10.48550/arXiv.1409.1556
- 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* (Salt Lake City, UT: IEEE), 1731–1740. doi: 10.1109/CVPR.2018.00186
- Su, Q., Chou, Y., Hu, Y., Li, J., Mei, S., Zhang, Z., et al. (2023). “Deep directly-trained spiking neural networks for object detection,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision* (Paris: IEEE), 6555–6565. doi: 10.1109/ICCV51070.2023.00603
- Sun, D. and Ji, H. (2023). “Event-based object detection using graph neural networks,” in *2023 IEEE 12th Data Driven Control and Learning Systems Conference (DDCLS)* (Xiangtan: IEEE), 1895–1900. doi: 10.1109/DDCLS58216.2023.10166491
- Sun, M., Zhang, D., Ge, Z., Wang, J., Li, J., Fang, Z., et al. (2024). Eventrpg: event data augmentation with relevance propagation guidance. *arXiv [Preprint]* arXiv:2403.09274. doi: 10.48550/arXiv.2403.09274
- Tan, M., and Le, Q. (2019). “Efficientnet: rethinking model scaling for convolutional neural networks,” in *International Conference on Machine Learning* (Long Beach, CA: PMLR), 6105–6114.
- Terven, J., Córdova-Esparza, D.-M., and Romero-González, J.-A. (2023). A comprehensive review of yolo architectures in computer vision: from yolov1 to yolov8 and yolo-nas. *Mach. Learn. Knowl. Extract.* 5, 1680–1716. doi: 10.3390/make5040083
- Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., et al. (2022). “Maxvit: multi-axis vision transformer,” in *European conference on computer vision* (Cham: Springer), 459–479. doi: 10.1007/978-3-031-20053-3\_27
- Wang, D., Jia, X., Zhang, Y., Zhang, X., Wang, Y., Zhang, Z., et al. (2023a). Dual memory aggregation network for event-based object detection with learnable representation. *Proc. AAAI Conf. Artif. Intell.* 37, 2492–2500. doi: 10.1609/aaai.v37i2.25346
- Wang, Y.-K., Wang, S.-E., and Wu, P.-H. (2023b). Spike-event object detection for neuromorphic vision. *IEEE Access* 11, 5215–5230. doi: 10.1109/ACCESS.2023.3236800
- Wang, Z., Wang, Z., Li, H., Qin, L., Jiang, R., Ma, D., et al. (2024). EAS-SNN: end-to-end adaptive sampling and representation for event-based detection with recurrent spiking neural networks. *arXiv [Preprint]*. arXiv:2403.12574. doi: 10.48550/arXiv.2403.12574
- Werbos, P. J. (1990). Backpropagation through time: what it does and how to do it. *Proc. IEEE* 78, 1550–1560. doi: 10.1109/5.58337
- Wu, Z., Gehrig, M., Lyu, Q., Liu, X., and Gilitschenski, I. (2024). “LEOD: label-efficient object detection for event cameras,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (Seattle, WA: IEEE), 16933–16943. doi: 10.1109/CVPR52733.2024.01602
- Xiao, X., Chen, X., Kang, Z., Guo, S., and Wang, L. (2022). “A spatio-temporal event data augmentation method for dynamic vision sensor,” in *International Conference on Neural Information Processing* (Cham: Springer), 422–433. doi: 10.1007/978-981-99-1645-0\_35
- Xu, G., Khan, A. S., Moshayedi, A. J., Zhang, X., and Shuxin, Y. (2022). The object detection, perspective and obstacles in robotic: a review. *EAI Endorsed Trans. AI Robot.* 1. doi: 10.4108/airo.v1i1.2709
- Yang, M., Liu, S.-C., Li, C., and Delbruck, T. (2012). “Addressable current reference array with 170db dynamic range,” in *2012 IEEE International Symposium on Circuits and Systems (ISCAS)* (Seoul: IEEE), 3110–3113. doi: 10.1109/ISCAS.2012.6271979
- Zhang, Y., Zhang, H., Li, Y., He, B., Fan, X., Wang, Y., et al. (2023). Direct training high-performance spiking neural networks for object recognition and detection. *Front. Neurosci.* 17:1229951. doi: 10.3389/fnins.2023.1229951
- Zhu, A. Z., Yuan, L., Chaney, K., and Daniilidis, K. (2018a). Ev-flownet: self-supervised optical flow estimation for event-based cameras. *arXiv [Preprint]*. arXiv:1802.06898. doi: 10.48550/arXiv.1802.06898
- Zhu, A. Z., Yuan, L., Chaney, K., and Daniilidis, K. (2018b). *Unsupervised event-based learning of optical flow, depth, and egomotion*.
- Zubić, N., Gehrig, D., Gehrig, M., and Scaramuzza, D. (2023). “From chaos comes order: Ordering event representations for object recognition and detection,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision* (Paris: IEEE), 12846–12856. doi: 10.1109/ICCV51070.2023.01180
- 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* (Seattle, WA: IEEE), 5819–5828. doi: 10.1109/CVPR52733.2024.00556