

CREST: An Efficient Conjointly-trained Spike-driven Framework for Event-based Object Detection Exploiting Spatiotemporal Dynamics

Ruixin Mao^{*}, Aoyu Shen^{*}, Lin Tang, Jun Zhou[†]

University of Electronic Science and Technology of China, Chengdu 611731, China
{mao.rx, shen_aoyu, tangl}@std.uestc.edu.cn, zhouj@uestc.edu.cn

Abstract

Event-based cameras feature high temporal resolution, wide dynamic range, and low power consumption, which are ideal for high-speed and low-light object detection. Spiking neural networks (SNNs) are promising for event-based object recognition and detection due to their spiking nature but lack efficient training methods, leading to gradient vanishing and high computational complexity, especially in deep SNNs. Additionally, existing SNN frameworks often fail to effectively handle multi-scale spatiotemporal features, leading to increased data redundancy and reduced accuracy. To address these issues, we propose CREST, a novel conjointly-trained spike-driven framework to exploit spatiotemporal dynamics in event-based object detection. We introduce the conjoint learning rule to accelerate SNN learning and alleviate gradient vanishing. It also supports dual operation modes for efficient and flexible implementation on different hardware types. Additionally, CREST features a fully spike-driven framework with a multi-scale spatiotemporal event integrator (MESTOR) and a spatiotemporal-IoU (ST-IoU) loss. Our approach achieves superior object recognition & detection performance and energy efficiency compared with state-of-the-art SNN algorithms on three datasets, providing an efficient solution for event-based object detection algorithms suitable for SNN hardware implementation.

Extended version — <http://arxiv.org/abs/2412.12525>

Introduction

Frame-based cameras excel in capturing photometric features but suffer from motion blur with fast-moving objects and are sensitive to extreme lighting conditions. Continuous full-frame processing also consumes high energy due to redundant background information. Conversely, event-based cameras capture changes in light intensity asynchronously at each pixel (Lichtsteiner, Posch, and Delbruck 2008; Gallego et al. 2020). This allows them to operate at much higher frequencies and under varied illumination conditions, ideal for high-speed and real-time object detection.

Deep ANN frameworks for event-based object detection encode event streams into dense, image-like representations,

leveraging conventional frame-based computer vision techniques (Peng et al. 2023; Zubic, Gehrig, and Scaramuzza 2024). Modern GPUs and TPUs also accelerate ANN training and inference (Nickolls et al. 2008; Jouppi et al. 2017). However, ANNs struggle with effectively handling the temporal features inherent in event-based data. This arises from losing crucial temporal dynamics and correlations when encoding events into discrete frames. Additionally, they consume high energy due to numerous MAC operations, which is problematic for energy-sensitive edge applications.

SNNs mimic the spiking feature of the biological neurons which are theoretically well-suited for processing event-based data (Maass 1997; Roy, Jaiswal, and Panda 2019). The timing and frequency of spikes convey diverse information, granting SNNs strong spatiotemporal characteristics. Moreover, SNNs perform accumulations (ACs) only when sparse spikes occur, making them inherently energy-efficient. However, SNNs often struggle with deep network structures due to inefficient training methods (Zhang et al. 2024; Nagaraj, Liyanagedera, and Roy 2023). Non-differentiable spikes require surrogate gradients for backpropagation (BP), potentially causing gradient vanishing in deep SNNs (Wu et al. 2018; Neftci, Mostafa, and Zenke 2019). The propagation of spatiotemporal gradients also greatly increases computation complexity. Some works convert well-trained ANN models to the same structured SNN. However, they require a large encoding time-window to approximate the ANN performance, causing a much higher spiking rate and energy consumption (Kim et al. 2020; Li et al. 2022b). More complex spiking neurons and layers with trainable parameters are proposed to lower the spiking rate, which add computation complexity for training & inference and are difficult to implement on existing SNN hardware (Fang et al. 2021; Cordone, Miramond, and Thierion 2022).

Furthermore, a unified SNN framework is needed to handle the multi-scale spatiotemporal features of event-based object detection. Existing approaches often fail to effectively encode the spatiotemporal characteristics of event data, leading to data redundancy and reduced accuracy (Cordone, Miramond, and Thierion 2022; Bodden et al. 2024). Moreover, the Intersection over Union (IoU) loss (Yu et al. 2016) employed in SNNs are generally adapted from ANNs, neglecting the unique spatiotemporal nature of event-based spike trains (Su et al. 2023; Fan et al. 2024). This may lead to inac-

^{*}These authors contributed equally.

[†]Corresponding author: Jun Zhou

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

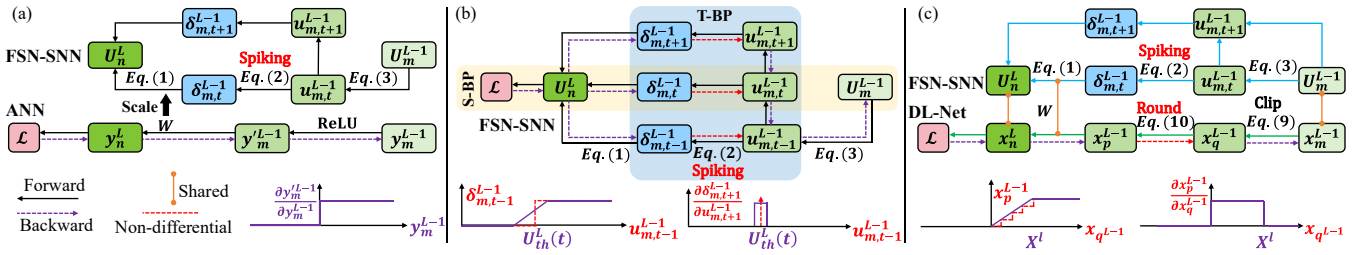


Figure 1: Comparison of (a) FSN-conversion, (b) FSN-BP, and (c) FSN-conjointly learning. The last layer uses the accumulated membrane potentials to make decisions. S-BP/T-BP denotes spatial-BP/temporal-BP.

curate regression and thus decrease the detection accuracy.

To this end, we propose *CREST*, an efficient conjointly-trained spike-driven framework exploiting spatiotemporal dynamics for event-based object detection. Firstly, we propose a simple yet effective conjoint learning rule with dual operation modes for efficient and flexible training implementation on different hardware types. For the backward process, we design a surrogate neural network with discrete-level activation values (DL-Net) to mimic the values represented by different spike train patterns. This replaces the original spatiotemporal gradient calculations, which reduces the computation complexity, alleviates the gradient vanishing problem, and speeds up the learning process compared with traditional SNN BP-like or conversion-based training.

Additionally, we propose a fully spike-driven framework for event-based object detection which includes a multi-scale spatiotemporal event integrator (MESTOR), a spatiotemporal-IoU (ST-IoU) loss, and few-spikes neuron (FSN) (Stöckl and Maass 2021) based SNN model. MESTOR not only aggregates event data across multiple scales but also extracts the spatiotemporal continuous events. This keeps the key spatiotemporal feature and reduces redundant background and noise events meantime. ST-IoU loss comprises the proposed spiking density-based IoU (Spiking-IoU) to exploit spatiotemporal continuity and the Complete-IoU (CIoU) (Zheng et al. 2021) for coordinate loss. The FSN-SNN adopts an efficient spike encoding scheme and is supported by the recently proposed high-performance SNN hardware STELLAR (Mao et al. 2024).

Our main contributions can be summarized as follows:

- ***CREST*, a spike-driven conjointly-trained framework exploiting the spatiotemporal dynamics** to enhance the efficiency of event-based object detection.
- **A novel conjoint learning rule** which introduces a surrogate neural network with discrete-level activation values to accelerate the learning process and alleviate the gradient vanishing issues in deep SNNs. Dual operation modes add flexibility to its hardware implementation.
- **A fully spike-driven framework** which incorporates the MESTOR, ST-IoU loss, and FSN-SNN model to handle multi-scale spatiotemporal features of the event-based object recognition and detection.
- Compared with the state-of-the-art SNN algorithms, our work achieves superior recognition & detection performance and up to 100× energy efficiency on 3 datasets.

Background and Motivation

Inefficiencies of LIF-like spiking neuron models. Leaky Integrate-and-Fire (LIF) neuron and its variants (details are shown in **Supplementary Material A**) are most commonly used in SNN algorithms and hardware implementations due to their trade-off between low computational complexity and biological interpretability (Abbott 1999; Gerstner and Kistler 2002; Fang et al. 2021). However, LIF neurons only support temporal or frequency encoding, which requires a long time window and many spikes to achieve high accuracy. Some introduce trainable parameters and normalization functions (e.g. complex exponent and division operations) to lower the spiking rate. However, the existing SNN hardware designs mainly adopt IF neurons for high energy efficiency which cannot support these complex operations.

Efficient Few-Spikes neuron (FSN). To improve encoding efficiency, the FSN and the FSN-based ANN-to-SNN conversion learning (FSN-conversion) method are proposed (Stöckl and Maass 2021). The spike trains encoded by FSNs carry both temporal and frequency information, which convey more information within a shorter time window than LIF-like neurons. Additionally, FSN-based SNN hardware demonstrates much higher energy efficiency and speedup compared to conventional SNN hardware (Mao et al. 2024). Therefore, we adopt FSNs as foundational units in this work.

The FSN works as follows. Firstly, the FSN n in layer l integrates the input spike δ_m^{l-1} to the membrane potential U_n^l (Eq. 1). N^{l-1} , K^{l-1} , $d^{l-1}(t)$, w_{mn}^l represent the number of input neurons, time-window length, spike coefficient at time t , and connection weights respectively. If $u_n^l(t)$ exceeds the threshold $U_{th}^l(t)$, a spike is fired (Eq. 2), and $u_n^l(t)$ decays by $U_{th}^l(t)$ (Eq. 3). The neuron dynamics parameters $U_{th}^l(t)$ and $d^l(t)$ can be used to model various activation functions.

$$U_n^l = \sum_{m=1}^{N^{l-1}} \sum_{t=1}^{K^{l-1}} d^{l-1}(t) \delta_m^{l-1}(t) w_{mn}^l \quad (1)$$

$$\delta_n^l(t) = \begin{cases} 1, & u_n^l(t) \geq U_{th}^l(t) \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

$$u_n^l(t+1) = U_n^l - \sum_t U_{th}^l(t) \delta_n^l(t) \quad (3)$$

Inefficiencies of the conventional SNN training methods. FSN-conversion (Fig. 1(a)): It uses neuron dynamics parameters to emulate the ANN activation functions, allowing the weights from a trained ANN to be directly applied to the

Algorithm 1: FS-Neuron based Conjoint Learning

Input: model weight w_{mn}^l from layer $l-1$ to l , input membrane potential U_m^{l-1}/x_m^{l-1} , threshold α^l and time-window K^l of layer l

Output: Updated model weight \mathbf{W}

```

1: Forward Pass:
2: if SNN – forward then //In SNN Hardware
3:   for  $t = 1$  to  $K^l$  do
4:      $\delta_m^{l-1}(t) \leftarrow \text{Spiking}(u_m^{l-1}(t), \alpha^{l-1}, K^{l-1})$     $\triangleright$  Eq.(2)
5:   end for
6:    $x_n^l = U_n^l \leftarrow \text{Integ.}(\delta_m^{l-1}(t), \alpha^{l-1}, K^{l-1}, w_{mn}^l)$     $\triangleright$  Eq.(1)
7: else //In GPU/TPU
8:    $x_q^{l-1} \leftarrow \text{Clip}(x_m^{l-1}, \alpha^{l-1}, K^{l-1})$     $\triangleright$  Eq.(9)
9:    $x_p^{l-1} \leftarrow \text{Round}(x_q^{l-1}, \alpha^{l-1}, K^{l-1})$     $\triangleright$  Eq.(10)
10:   $x_n^l \leftarrow \text{Convolution}(x_p^{l-1}, w_{mn}^l)$     $\triangleright$  Eq.(11)
11: end if
12: Backward Pass:
13: //Surrogate gradient BP in DL-Net
14:  $\frac{\partial x_n^l}{\partial x_m^{l-1}} = \frac{\partial x_n^l}{\partial x_p^{l-1}} \frac{\partial x_p^{l-1}}{\partial x_q^{l-1}} \approx \frac{\partial x_n^l}{\partial x_p^{l-1}}$ 
15: return Updated model weight  $\mathbf{W}$ 

```

same-structured SNN. Despite its better performance than LIF-based SNNs, it still limits the accuracy with the short time-window due to the difference between ANN and FSN functions. Increasing the time-window improves the approximation between FSN and ANN functions thus reducing the conversion accuracy loss. But this also increases spiking rate and energy consumption as demonstrated in (Mao et al. 2024). FSN based BP (FSN-BP) learning (Fig. 1(b)): Similar to the traditional SNN BP-like learning, FSN-BP can be achieved with a surrogate gradient for non-differentiable spike. The spatiotemporal backpropagation is denoted as Eq. 4, with ϵ_m^{L-1} as error from loss \mathcal{L} to layer $L-1$. The surrogate gradient is approximated as Eq. 5, where rect equals 1 when the input is within $[-0.5, 0.5]$ and 0 otherwise. This approach achieves higher accuracy with a shorter time-window and fewer spikes than FSN-Conversion. However, it has high computation complexity due to the spatiotemporal gradient calculation. Further, deep networks exacerbate the issue of handling and storing sparse spike timing information. Gradients become very small and even diminish during backpropagation, pronouncing the gradient vanishing problem.

$$\epsilon_m^{L-1} = \frac{\partial \mathcal{L}}{\partial U_m^{L-1}} = \sum_{n=1}^{n^L} \sum_{t=1}^{K^L} w_{mn}^L \frac{\partial \delta_m^{L-1}(t)}{\partial u_m^{L-1}(t)} d^L(t) \epsilon_n^L \quad (4)$$

$$\frac{\partial \delta_m^{L-1}(t)}{\partial u_m^{L-1}(t)} = \frac{1}{2\beta} \cdot \text{rect} \left(\frac{u_m^{L-1}(t) - U_{th}^{L-1}(t)}{2\beta} \right) \quad (5)$$

Method

We propose *CREST*, a spike-driven conjointly-trained framework exploiting the spatiotemporal dynamics. We establish a dual-model conjoint learning rule to simplify the spatiotemporal BP which support efficient implementation on different hardware types. Based on this, a fully spike-driven framework is developed to handle multi-scale spatiotemporal features of event-based object detection.

FS-Neuron based Conjoint Learning Rule

We first demonstrate how to use FSN to emulate ReLU to give an insight into the FSN computation mechanism. Then, a surrogate DL-Net is proposed to mimic the continuous values represented by discrete spike train patterns. Further, we introduce the spatial surrogate gradient backpropagation to DL-Net. Finally, we illustrate the operation flow of the conjoint learning rule.

Similarity between FSN and ReLU Functions. ANN-based object detection mostly adopts ReLU and leaky-ReLU (Glorot, Bordes, and Bengio 2011; Maas et al. 2013). We implement ReLU in this paper and the detailed reasons are illustrated in **Supplementary Material B**. Based on Eq. 3, rewrite the membrane potential U_m^{l-1} as Eq. 6. Meanwhile, denote the unweighted integration of output spikes from FSN m as Eq. 7 based on Eq. 1.

$$U_m^{l-1} = \sum_{t=1}^{K^{l-1}} U_{th}^{l-1}(t) \delta_m^{l-1}(t) + u_m^{l-1}(K^{l-1}) \quad (6)$$

$$U_p^{l-1} = \sum_{t=1}^{K^{l-1}} d^{l-1}(t) \delta_m^{l-1}(t) \quad (7)$$

With the help of FSN parameters, the membrane potential is encoded into a spike train and then decoded into a continuous value. In this process, U_m^{l-1} and U_p^{l-1} emulate the input and output of ReLU, and the MAC in ANNs are replaced by the ACs of the corresponding weight at spiking time steps in SNNs. To preserve non-linearity while controlling the output range, ReLU is bounded by α . This is formulated as three cases: **1** when $U_m^{l-1} \leq 0$, $U_p^{l-1} = 0$; **2** when $0 < U_m^{l-1} \leq \alpha$, $U_p^{l-1} = U_m^{l-1}$; and **3** when $U_m^{l-1} > \alpha$, $U_p^{l-1} = \alpha$. Case **2** is optimized by minimizing the error with proper FSN parameters under the distribution of U_m^{l-1} (Eq. 8), with K^{l-1} pre-set by the requirements of accuracy and firing rate.

$$\text{min } \varepsilon(U_{th}^{l-1}(t), d^{l-1}(t); K^{l-1}) = \mathbb{E}_{u_m(t)} |U_m^{l-1} - U_p^{l-1}| \quad (8)$$

A possible setting is $U_{th}^{l-1}(t) = d^{l-1}(t) = \alpha^{l-1} 2^{-t}$. For case **1**, when $U_m^{l-1} \leq 0$, FSN m generates no spikes ($U_p^{l-1} = 0$). For case **2**, when U_m^{l-1} is an integer multiple of $\alpha^{l-1} 2^{-K^{l-1}}$, $U_p^{l-1} = U_m^{l-1}$. Otherwise, U_m^{l-1} rounds down to the nearest level with a tolerance $\alpha^{l-1} 2^{-K^{l-1}}$. For case **3**, $U_m^{l-1}(t) \geq \alpha^{l-1} 2^{-t}$, the spike train is encoded as K -bit 1 with $U_p^{l-1} = \alpha^{l-1} (1 - 2^{-K^{l-1}})$. The error between U_p^{l-1} and U_m^{l-1} decreases when K^{l-1} increases. Noted that both the earlier spikes and the higher spiking frequency represent larger values (more significant information), as reflected by larger coefficients ($d(t)$) and accumulated values, which improves encoding efficiency.

Surrogate DL-Net. FSN-BP requires complex and iterative spatiotemporal gradient BP with high computation complexity, which can exacerbate gradient vanishing in deep SNNs. Inspired by the coarse-to-fire-processing strategy of FSN-ReLU, we design a surrogate neural network with discrete-level activation values (DL-Net) to approximate spike train patterns in Fig. 1(c).

To emulate the boundary condition of the 3 cases, x_m^{l-1} is clipped to x_q^{l-1} (Eq. 9) with $X^{l-1} = \alpha^{l-1} (1 - 2^{-K^{l-1}})$. Fur-

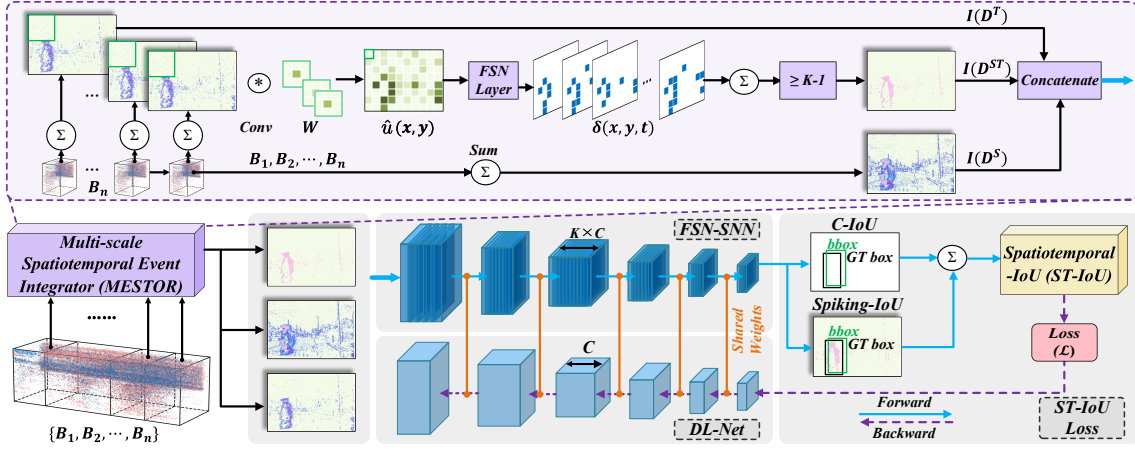


Figure 2: The overview architecture of CREST.

ther, x_q^{l-1} is discretized to mimic the FSN spiking process with $X_{\min}^{l-1} = \alpha^{l-1} 2^{-K^{l-1}}$ (Eq. 10). Finally, the temporal iterative ACs in Eq. 1 are replaced by MACs (Eq. 11), where x_n^l emulates U_n^l . The connecting weights w_{pn}^l in DL-Net are shared with those in the FSN-SNN.

$$x_q^{l-1} = \text{clip}(x_m^{l-1}) = \begin{cases} 0, & x_m^{l-1} < 0 \\ x_m^{l-1}, & 0 \leq x_m^{l-1} \leq X^{l-1} \\ X^{l-1}, & \text{otherwise} \end{cases} \quad (9)$$

$$x_p^{l-1} = \text{round}\left(\frac{x_q^{l-1}}{X_{\min}^{l-1}}\right) \times X_{\min}^{l-1} \quad (10)$$

$$x_n^l = \sum_{p=1}^{N^{l-1}} x_p^{l-1} w_{pn}^l \quad (11)$$

Spatial surrogate gradient backpropagation in DL-Net
The round function in Eq. 10 causes abrupt changes, leading to a discontinuous derivative. We replace the round function with a smooth identity function during backpropagation as shown in the bottom of Fig.1 (c). Therefore, the iterative surrogate gradient calculation for spatiotemporal spiking in Eq. 4 is replaced by the surrogate gradient (Eq. 12) for spatial discretization, which greatly reduces the computation complexity and improves training efficiency. Additionally, the sparse spikes are aggregated into a continuous value, which avoids small and unstable gradients, greatly alleviating the gradient vanishing problem. The detailed process of BP in DL-Net is illustrated in **Supplementary Material C**.

$$\frac{\partial x_p^l}{\partial x_q^l} \approx 1 \quad (12)$$

Dual training modes for flexible hardware implementation (Algorithm 1). Case A: GPU/TPU implementation. GPUs/TPUs primarily focus on accelerating parallel matrix operations, prioritizing speed over energy efficiency. Thus, DL-Net is used for forward & backward passes in training (green & purple line in Fig.1 (c)). Case B: SNN Hardware. FSN-based SNN hardware is specifically designed to accelerate FSN computation with ultra-low energy consumption. Therefore, FS-SNN/DL-Net can be used for the for-

ward/backward pass (blue/purple line), which share membrane potentials and weights (orange line).

Overview of Spike-driven Framework for Event-based Object Detection

Fig. 2 shows the overall architecture of the spike-driven framework for event-based object detection which includes MESTOR, ST-IoU loss, and FSN-SNN model. MESTOR integrates the input feature from multi-scales. The integrated features are then sent to the FSN-based SNN structure. The FSN model is chosen based on its high encoding efficiency, efficient hardware implementation (Mao et al. 2024), and the proposed efficient conjointly training method mentioned above. Additionally, the ST-IoU loss is proposed to improve detection accuracy for the spatiotemporal event-based data.

Multi-scale Spatiotemporal Event Integrator

For a continuous event-based period I , the raw stream $E = e_i(x_i, y_i, t_i, p_i)_{i \in M}$ with total M bipolar events p_i at pixel point (x_i, y_i) exhibits high sparsity and temporal redundancy. These characteristics impede feature learning and increase computational overhead. Moreover, the raw stream contains many redundant background and noise events, which can disrupt the learning process. Existing methods mostly split $I(E)$ into equal bins and integrate the events in each bin to obtain a spatiotemporal dense representation (Eq. 13) (Su et al. 2023; Fan et al. 2024). The bin length affects the spatiotemporal feature preservation, leading to a trade-off between both scales. We aim to integrate the input feature from multi-scales (spatiotemporal scale \mathcal{ST} , spatial scale \mathcal{S} and temporal scale \mathcal{T}) for better representation.

$$I(E) \Rightarrow I(D^{ST}) \quad (13)$$

Spatiotemporal continuous $I(D^{ST})$. We observe that the events generated by the same object exhibit both time and space continuity, whereas background and noise events are less time-continuous. Therefore, we leverage the spatiotemporal properties of the spiking neuron to extract and cluster the spatiotemporal continuous events and reduce redundant background and noise events. Like (Nagaraj, Liyan-

Method	Representation	Net	Params	Acc	K	fr/spar	AC/MAC	Energy(mJ)
YOLO (Cannici et al. 2019)	VoxelGrid	ANN	-	0.927	-	1	0.16G	0.75
EvS-S (Li et al. 2021)	Graph	GNN	-	0.931	-	-	-	-
Asynet (Messikommer et al. 2020)	VoxelGrid	ANN	-	0.944	-	0.067	0.16G	0.05
HybridSNN (Kugele et al. 2021)	HIST	SNN	-	0.770	-	-	-	-
Gabor-SNN (Sironi et al. 2018)	HAT	SNN	-	0.789	-	-	-	-
Squeeze-1.1 (Cordone, Miramond, and Thierion 2022)	VoxelCube	SNN	0.72M	0.846	5	0.251	0.02G	0.36
Mobile-64 (Cordone, Miramond, and Thierion 2022)	VoxelCube	SNN	18.81M	0.917	5	0.171	4.20G	83.34
Dense121-24 (Cordone, Miramond, and Thierion 2022)	VoxelCube	SNN	3.93M	0.904	5	0.336	2.25G	37.76
VGG-11 (Cordone, Miramond, and Thierion 2022)	VoxelCube	SNN	9.23M	0.924	5	0.120	0.61G	12.68
DenseNet121-16 (Fan et al. 2024)	VoxelCube	SNN	1.76M	0.937	5	0.147	0.06G*	1.22
<i>CREST</i> (CSPdarknet-tiny)	MESTOR	SNN	3.41M	0.949	5	0.165	0.05G	0.04
<i>CREST</i> (DenseNet121-16)	MESTOR	SNN	1.95M	0.952	5	0.146	0.07G	0.04
<i>CREST</i> (ShuffleNetV2)	MESTOR	SNN	1.01M	0.940	5	0.176	0.01G	0.01

* denotes self-implementation result; Acc denotes Accuracy.

Table 1: Comparison between *CREST* and the state-of-the-art methods on NCARs dataset.

agedera, and Roy 2023), we first divide $I(E)$ into N equal-length (Δt) time bins $\mathcal{B}(x, y, n)$ without polarity.

$$\mathcal{B}(x, y, n) = \sum_{i=1}^{(n-1)\Delta t < e(t_i) < n\Delta t} e_i(x_i, y_i, t_i) \quad (14)$$

An FSN-based convolution layer is constructed which applies fixed-value filters of size 3×3 to time bin \mathcal{B}_t . In detail, the convolution results of each time bin and each filter are accumulated to the membrane potential \hat{u} in each time step (Eq. 15). Then, based on Eq.2 and 3, each neuron generates a spike train (Eq. 16). With a given time window K , if the number of spikes in one spike train is greater than $K - 1$, this neuron (pixel) is considered space-time continuous and we set the value of this pixel as 1 (we retain this pixel). Otherwise, we abandon this pixel (Eq. 17).

$$\hat{u}(x, y, t) = \mathcal{B}_t * \mathbf{W} + u(x, y, t) \quad (15)$$

$$\{\delta(x, y, t), u(x, y, t + 1)\} = \text{FSN}(\hat{u}(x, y, t)) \quad (16)$$

$$I(D^{S\mathcal{T}})(x, y) = \begin{cases} 1, & \sum_{t=1}^K \delta(x, y, t) \geq K - 1 \\ 0, & \text{otherwise} \end{cases} \quad (17)$$

Spatial scale $I(D^S)$ and temporal scale $I(D^{\mathcal{T}})$. Combined with Eq. 14, we observe that shorter time bins capture finer temporal details but increase data sparsity, while longer intervals enhance spatial feature density but may smooth out crucial temporal features. So we use short time bins to keep temporal features (Eq. 18) and use the accumulation of short time bins to keep spatial features (Eq. 19).

$$I(D^{\mathcal{T}})(x, y) = \mathcal{B}(x, y, N) \quad (18)$$

$$I(D^S)(x, y) = \sum_{n=1}^N \mathcal{B}(x, y, n) \quad (19)$$

Finally, the features extracted from multi-scales are integrated as a 3-channel input and fed into the following network as Eq. (20).

$$I(E) \implies I(D^{S\mathcal{T}}) \& I(D^S) \& I(D^{\mathcal{T}}) \quad (20)$$

Spatiotemporal IoU Loss

Bounding box regression loss plays a key role in object detection. IoU, GIoU (Rezatofighi et al. 2019), and CIoU (Zheng et al. 2021) are widely used in frame-based object

detection and we show the loss function of YOLOv4 in Eq. 21, where \mathcal{L}_{obj} is confidence loss, \mathcal{L}_{cls} is classification loss and \mathcal{L}_{reg} is bounding box regression loss. Current event-based detection methods often directly apply these frame-based approaches, which compute the coordinate differences between the predicted boxes and gtboxes. This neglects the unique spatiotemporal characteristics of event-based spike trains, potentially leading to less precise regression.

$$\begin{aligned} \mathcal{L}_{total} &= \mathcal{L}_{obj} + \mathcal{L}_{cls} + \mathcal{L}_{reg} \\ \mathcal{L}_{reg} &= 1 - CIoU \end{aligned} \quad (21)$$

As previously noted, the events generated by the target objects exhibit spatiotemporal continuity. This can be represented through the spike density ρ (total number of spikes within a box divided by its area). ρ around the object is usually higher than around the background. We propose Spiking-IoU to quantify the difference in spike density between the predicted $B = (x, y, w, h)$ and ground truth $B^{gt} = (x^{gt}, y^{gt}, w^{gt}, h^{gt})$ boxes (Eq. 22).

$$\text{Spiking-IoU} = |\rho_{sd}^{gt} - \rho_{sd}| = \left| \frac{\sum_{m=1}^M \delta_m^{gt}}{w^{gt} h^{gt}} - \frac{\sum_{n=1}^N \delta_n}{wh} \right| \quad (22)$$

Furthermore, we propose ST-IoU loss to fully exploit the spatiotemporal feature in the event-based detection (Eq. 23). It incorporates the conventional IoU scheme to provide accurate coordinate information for fast convergence and Spiking-IoU to enhance detection precision with unique spatiotemporal information. a and b are constant weights.

$$\begin{aligned} \text{ST-IoU} &= a \times \text{Spiking-IoU} + b \times \text{CIoU} \\ \mathcal{L}_{reg} &= 1 - \text{ST-IoU} \end{aligned} \quad (23)$$

Experiments

We first describe the experimental settings. Then, the effective tests and ablation tests are conducted to verify the effectiveness and performance of the proposed methods. Additionally, we analyze *CREST* in exploiting sparsity.

Experimental Settings

Datasets. We adopt the commonly used NCARs dataset for object recognition (Sironi et al. 2018), Gen1 and PKU-Vidar-DVS dataset for object detection (de Tournemire et al. 2020; Li et al. 2022a). The details of these datasets are

Method	Rep.	Net	Head	Params	mAP _{50:95}	mAP ₅₀	K	fr/sp	AC/MAC	Energy
Asynet (Messikommer et al. 2020)	VoxelGrid	ANN	YOLO	11.4M	0.129	-	-	0.100	1.05G	0.48
S-Center (Bodden et al. 2024)	HIST	ANN	CenterNet	12.97M	0.278	-	-	1	6.13G	28.21
EGO-12 (Zubić et al. 2023)	EGO-12	ANN	YOLOv6	140M*	0.504	-	-	1	84.34G*	387.96
VC-Dense (Cordone, Miramond, and Thierion 2022)	VoxelCube	SNN	SSD	8.2M	0.189	-	5	0.372	2.33G	37.55
VC-Mobile (Cordone, Miramond, and Thierion 2022)	VoxelCube	SNN	SSD	24.26M	0.147	-	5	0.294	4.34G	76.22
LT-SNN (Hassan, Meng, and sun Seo 2023)	HIST	SNN	YOLOv2	86.82M*	0.298	-	-	-	19.53G*	-
S-Center (Bodden et al. 2024)	HIST	SNN	CenterNet	12.97M	0.229	-	5	0.174	6.38G	126.21
EMS-10 (Su et al. 2023)	HIST	SNN	YOLOv3	6.20M	0.267	0.547	5	0.211	5.90G*	112.83
EMS-18 (Su et al. 2023)	HIST	SNN	YOLOv3	9.34M	0.286	0.565	5	0.201	9.70G*	187.02
EMS-34 (Su et al. 2023)	HIST	SNN	YOLOv3	14.40M	0.310	0.590	5	0.178	32.99G*	650.13
TR-YOLO (Yuan et al. 2024)	HIST	SNN	YOLOv3	8.7M	-	0.451	3	-	-	-
SFOD (Fan et al. 2024)	VoxelCube	SNN	SSD	11.9M	0.321	0.593	5	0.240	6.72G*	124.73
CREST (CSPdarknet-tiny)		SNN	YOLOv4	7.61M	0.360	0.632	5	0.167	8.39G	6.31
CREST (DenseNet121-16)	MESTOR	SNN	YOLOv4	3.01M	0.339	0.615	5	0.095	8.15G	3.48
CREST (ShuffleNetV2)		SNN	YOLOv4	1.79M	0.305	0.568	5	0.208	2.42G	2.27

* denotes self-implementation results; Rep. denotes Representation; energy unit is mJ.

Table 2: Comparison between *CREST* and the state-of-the-art methods on Gen1 dataset.




Figure 3: Gen1/PKU-Vidar-DVS encoded with MESTOR.

shown in **Supplementary Material D**.

Implementation Details. Set $K = 5$ and $\alpha = 3$ for middle layers, with $\alpha = 1$ for the input encoding layer in FSN-SNN models. All models are trained on a Tesla V100 GPU using the AdamW optimizer and StepLR scheduler. For recognition, training spans 100 epochs with a batch size of 64, an initial learning rate of $1e - 3$, and a weight decay of $1e - 2$. For detection, training also spans 100 epochs with a batch size of 32, an initial learning rate of $1e - 3$, and a weight decay of $5e - 4$. Data augmentation techniques include mosaic, random rotation, crop and flipping.

Performance Metrics. Accuracy & mAP_{50:95} & mAP₅₀ are used for object recognition and detection respectively. We also establish the **energy consumption model** to show the efficiency of *CREST*. Present works only consider the energy consumed by spike integration, which neglect the energy required for the decay operation of LIF-like neurons despite its inescapability in hardware implementation (Fan et al. 2024; Su et al. 2023; Dampfhofer et al. 2022).

For ANNs, energy consumption is assessed by the number of MACs (OP_{MAC}) (Eq. 24). For fairness, we also account for sparsity (sp) in methods that leverage spatiotemporal sparsity (Messikommer et al. 2020). For LIF-based SNNs, energy consumption includes spike integration and extra MACs caused by membrane poten-

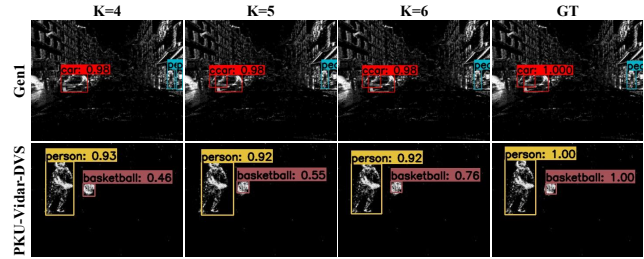


Figure 4: Visualized detection results with different K.

tial decay (Eq. 25). For FSN-based SNNs, as mentioned in Background and Motivation, they have negligible decay energy costs (Eq. 26). The energy cost for per 32-bit floating-point AC/MAC operation is 0.9/4.6 pJ (Horowitz 2014).

$$E_{ANN} = OP_{MAC} \times sp \times E_{MAC} \quad (24)$$

$$E_{LIF-SNN} = K \times OP_{AC} \times fr \times E_{AC} + K \times OP_{AC} \times (1 - fr) \times E_{MAC} \quad (25)$$

$$E_{FS-SNN} = K \times OP_{AC} \times fr \times E_{AC} \quad (26)$$

Effective Tests

Object Recognition. NCAR samples are resized to 64×64 pixels, encoded with MESTOR and conjointly-trained with three FSN-based networks (Bochkovskiy, Wang, and Liao 2020; Huang et al. 2017; Ma et al. 2018). Our model outperforms existing methods in accuracy with smaller networks in Table 1. Moreover, its lower firing rate reduces energy consumption compared to sparse ANNs and SOTA SNNs.

Object Detection. We evaluate *CREST* on the Gen1 and PKU-Vidar-DVS (9 classes under extreme lighting conditions for event-based object detection) datasets. The visualization of the MESTOR encoding process and the detection results are shown in Fig. 3 and Fig. 4. As shown in Table 2, *CREST* achieves high performance with fewer parameters and ultra-low energy consumption. It improves mAP_{50} from 0.59 to 0.63 compared to SOTA SNNs, while reducing energy consumption by 100×. Even our smallest network

Representation	Backbone	K	fr/sp	mAP ₅₀
HIST (Chen 2018)	YOLO ANN	1	0.331	
HIST (Iacono et al. 2018)	SSD ANN	1	0.326	
EV. (Hu, Delbruck, and Liu 2020)	YOLOv3 ANN	1	0.353	
TAR-events (Li et al. 2022a)	YOLOv3 ANN	1	0.386	
MESTOR	CSPdark	5	0.173	0.432
	Dense121	5	0.092	0.429
	ShuffleV2	5	0.195	0.425

EV. denotes EventVolume.

Table 3: Comparison between *CREST* and the state-of-the-art methods on PKU-DVS-Vidar dataset.

Method	CLR	MESTOR	ST-IOU	mAP _{50:95}	fr	Time
Baseline				0.269	0.184	1×
CREST-A	✓			0.279	0.185	3.42×
CREST-B	✓	✓		0.358	0.173	3.43×
<i>CREST</i>	✓	✓	✓	0.360	0.167	3.41×

Table 4: Ablation study: comparison between the baseline, CREST-A, CREST-B, and *CREST* on Gen1.

(1.79M parameters) performs competitively with existing SNNs, whereas the sparse ANN (Messikommer et al. 2020) lag behind in mAP. Note that the highest mAP in (Zubić et al. 2023) is achieved with a much larger network. In Table 3, *CREST* performs well under these challenging conditions, surpassing existing ANN methods in mAP.

Ablation Studies

Overview. We conducted ablation studies to validate the effectiveness of each component in *CREST* (Table 4). The baseline is trained with FSN-BP, $K = 5$, CSPdarknet-tiny backbone, YOLOv4-head, CIOU regression loss, and histogram representation. CREST-A replaces FSN-BP in the baseline with CLR, and CREST-B further substitutes the representation in CREST-A with MESTOR.

The conjoint learning rule reduces training time per epoch by 3.4× and enhances mAP performance by reducing errors caused by surrogate gradient backpropagation across multiple dimensions. MESTOR and ST-IOU further boost mAP and lower firing rates by fully leveraging the spatiotemporal feature of the event data.

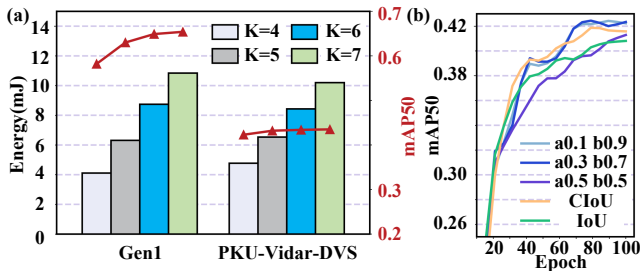


Figure 5: (a) Energy and mAP change with time-window on CSPdarknet-tiny ; (b) mAP of Shufflenet on PKU-Vidar-DVS with different IoUs.

Method	Acc		
	SFOD	w/ FSN-BP(ours)	w/ CLR(ours)
DenseNet121-16	0.937	0.948	0.952
DenseNet121-32	0.923	0.948	0.953
DenseNet169-16	0.921	0.946	0.950
DenseNet169-32	0.894	0.944	0.953

Table 5: Ablation study: comparison between the conventional learning rule and the proposed CLR on NCARs.

Representation	CSPdark-tiny		Dense121-16		ShuffleV2	
	Acc	fr	Acc	fr	Acc	fr
MESTOR-A	0.873	0.164	0.871	0.151	0.859	0.187
MESTOR-B	0.938	0.185	0.937	0.154	0.922	0.199
MESTOR	0.949	0.165	0.952	0.146	0.940	0.176

Table 6: Ablation study: effectiveness of MESTOR’s components on NCARs dataset.

Conjoint Learning Rule. We validate CLR’s effectiveness in mitigating gradient vanishing and enhancing performance in deep SNNs (Table 5). LIF-based SFOD trained with STBP (Wu et al. 2018) struggle with deep SNNs, resulting in low accuracy. In contrast, our model trained with CLR demonstrates the accuracy improvement compared with both SFOD and our model trained with FSN-BP.

Multi-scale Spatiotemporal Event Integrator. We denote spatiotemporal continuous $\mathcal{I}(D^{ST})$ as A, and spatial $\mathcal{I}(D^S)$ & temporal $\mathcal{I}(D^T)$ as B. Ablation experiments on MESTOR components (Table 6) show that A reduces the firing rate by reducing redundant events through extracting and clustering the spatiotemporal continuous events, while B improves accuracy by retaining key spatiotemporal features. Integrating A and B (MESTOR) achieves the highest accuracy and a comparatively low firing rate.

Spatiotemporal-IoU. Fig. 5(b) shows *CREST* with different IoUs. ST-IoU achieves higher mAP than CIOU. Experiments with different a&b in Eq. 23 show that in ST-IoU, the CIOU part accelerates the convergence speed, while the Spiking-IoU part improves the mAP due to its ability to fully exploit the spatiotemporal feature of the event-based data.

Additional analysis of *CREST* in exploiting sparsity. Fig. 5(a) shows the energy consumption and accuracy of CSPdarknet-tiny on two datasets with varying degrees of sparsity (i.e. firing rates under different time windows). The energy cost is reduced by up to 53.2% with nearly unpruned mAP when the firing rate decays.

Conclusion

We propose *CREST*, an innovative conjointly-trained spike-driven framework (incorporate MESTOR, ST-IoU loss, and FSN-SNN model) tailored for high accuracy and energy-efficient event-based object detection. With up to 100× improvement in energy efficiency over SOTA SNN algorithms, *CREST* offers a groundbreaking solution for advanced event-based object detection systems.

Acknowledgments

This work was supported by NSAF under Grant U2030204. (Corresponding author: Jun Zhou.)

References

- Abbott, L. F. 1999. Lapicque's introduction of the integrate-and-fire model neuron (1907). *Brain research bulletin*, 50(5-6): 303–304.
- Bochkovskiy, A.; Wang, C.-Y.; and Liao, H.-Y. M. 2020. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*.
- Bodden, L.; Schwaiger, F.; Ha, D. B.; Kreuzberg, L.; and Behnke, S. 2024. Spiking CenterNet: A Distillation-boosted Spiking Neural Network for Object Detection. In *International Joint Conference on Neural Networks (IJCNN)*.
- Cannici, M.; Ciccone, M.; Romanoni, A.; and Matteucci, M. 2019. Asynchronous convolutional networks for object detection in neuromorphic cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*.
- Chen, N. F. 2018. Pseudo-labels for supervised learning on dynamic vision sensor data, applied to object detection under ego-motion. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 644–653.
- 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)*, 1–8. IEEE.
- Dampfhofer, M.; Mesquida, T.; Valentian, A.; and Anghel, L. 2022. Are SNNs really more energy-efficient than ANNs An in-depth hardware-aware study. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 7(3): 731–741.
- de Tournemire, P.; Nitti, D.; Perot, E.; Migliore, D.; and Sironi, A. 2020. A Large Scale Event-based Detection Dataset for Automotive. *arXiv:2001.08499*.
- Fan, Y.; Zhang, W.; Liu, C.; Li, M.; and Lu, W. 2024. SFOD: Spiking Fusion Object Detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17191–17200.
- Fang, W.; Yu, Z.; Chen, Y.; Masquelier, T.; Huang, T.; and Tian, Y. 2021. Incorporating learnable membrane time constant to enhance learning of spiking neural networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2661–2671.
- Gallego, G.; Delbrück, T.; Orchard, G.; Bartolozzi, C.; Taba, B.; Censi, A.; Leutenegger, S.; Davison, A. J.; Conradt, J.; Daniilidis, K.; et al. 2020. Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(1): 154–180.
- Gerstner, W.; and Kistler, W. M. 2002. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press.
- Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 315–323. JMLR Workshop and Conference Proceedings.
- Hasssan, A.; Meng, J.; and sun Seo, J. 2023. LT-SNN: Self-Adaptive Spiking Neural Network for Event-based Classification and Object Detection.
- Horowitz, M. 2014. 1.1 computing's energy problem (and what we can do about it). In *2014 IEEE international solid-state circuits conference digest of technical papers (ISSCC)*, 10–14. IEEE.
- Hu, Y.; Delbruck, T.; and Liu, S.-C. 2020. Learning to exploit multiple vision modalities by using grafted networks. In *European Conference on Computer Vision*, 85–101. Springer.
- 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*, 4700–4708.
- Iacono, M.; Weber, S.; Glover, A.; and Bartolozzi, C. 2018. Towards event-driven object detection with off-the-shelf deep learning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1–9. IEEE.
- Jouppi, N. P.; Young, C.; Patil, N.; Patterson, D.; Agrawal, G.; Bajwa, R.; Bates, S.; Bhatia, S.; Boden, N.; Borchers, A.; et al. 2017. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th annual international symposium on computer architecture*, 1–12.
- Kim, S.; Park, S.; Na, B.; and Yoon, S. 2020. Spiking-yolo: spiking neural network for energy-efficient object detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 11270–11277.
- 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*, 297–312. Springer.
- Li, J.; Wang, X.; Zhu, L.; Li, J.; Huang, T.; and Tian, Y. 2022a. Retinomorph object detection in asynchronous visual streams. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 1332–1340.
- Li, Y.; He, X.; Dong, Y.; Kong, Q.; and Zeng, Y. 2022b. Spike calibration: Fast and accurate conversion of spiking neural network for object detection and segmentation. *arXiv preprint arXiv:2207.02702*.
- Li, Y.; Zhou, H.; Yang, B.; Zhang, Y.; Cui, Z.; Bao, H.; and Zhang, G. 2021. Graph-based asynchronous event processing for rapid object recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 934–943.
- Lichtsteiner, P.; Posch, C.; and Delbruck, T. 2008. A 128 × 128 120 dB 15 μs latency asynchronous temporal contrast vision sensor. *IEEE journal of solid-state circuits*, 43(2): 566–576.
- Ma, N.; Zhang, X.; Zheng, H.-T.; and Sun, J. 2018. Shufflenet v2: Practical guidelines for efficient cnn architecture

- design. In *Proceedings of the European conference on computer vision (ECCV)*, 116–131.
- Maas, A. L.; Hannun, A. Y.; Ng, A. Y.; et al. 2013. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, 3. Atlanta, GA.
- Maass, W. 1997. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9): 1659–1671.
- Mao, R.; Tang, L.; Yuan, X.; Liu, Y.; and Zhou, J. 2024. Stellar: Energy-Efficient and Low-Latency SNN Algorithm and Hardware Co-Design with Spatiotemporal Computation. In *2024 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, 172–185. IEEE.
- 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*, 415–431. Springer.
- Nagaraj, M.; Liyanagedera, C. M.; and Roy, K. 2023. Dotie-detecting objects through temporal isolation of events using a spiking architecture. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 4858–4864. IEEE.
- Neftci, E. O.; Mostafa, H.; and Zenke, F. 2019. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6): 51–63.
- Nickolls, J.; Buck, I.; Garland, M.; and Skadron, K. 2008. Scalable parallel programming with cuda: Is cuda the parallel programming model that application developers have been waiting for? *Queue*, 6(2): 40–53.
- Peng, Y.; Zhang, Y.; Xiao, P.; Sun, X.; and Wu, F. 2023. Better and faster: Adaptive event conversion for event-based object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 2056–2064.
- Rezatofighi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I.; and Savarese, S. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 658–666.
- Roy, K.; Jaiswal, A.; and Panda, P. 2019. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784): 607–617.
- 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*, 1731–1740.
- Stöckl, C.; and Maass, W. 2021. Optimized spiking neurons can classify images with high accuracy through temporal coding with two spikes. *Nature Machine Intelligence*, 3(3): 230–238.
- Su, Q.; Chou, Y.; Hu, Y.; Li, J.; Mei, S.; Zhang, Z.; and Li, G. 2023. Deep directly-trained spiking neural networks for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6555–6565.
- Wu, Y.; Deng, L.; Li, G.; Zhu, J.; and Shi, L. 2018. Spatio-temporal backpropagation for training high-performance spiking neural networks. *Frontiers in neuroscience*, 12: 331.
- Yu, J.; Jiang, Y.; Wang, Z.; Cao, Z.; and Huang, T. 2016. Unitbox: An advanced object detection network. In *Proceedings of the 24th ACM international conference on Multimedia*, 516–520.
- Yuan, M.; Zhang, C.; Wang, Z.; Liu, H.; Pan, G.; and Tang, H. 2024. Trainable Spiking-YOLO for low-latency and high-performance object detection. *Neural Networks*, 172: 106092.
- Zhang, H.; Li, Y.; Leng, L.; Che, K.; Liu, Q.; Guo, Q.; Liao, J.; and Cheng, R. 2024. Automotive Object Detection via Learning Sparse Events by Spiking Neurons. *IEEE Transactions on Cognitive and Developmental Systems*.
- Zheng, Z.; Wang, P.; Ren, D.; Liu, W.; Ye, R.; Hu, Q.; and Zuo, W. 2021. Enhancing geometric factors in model learning and inference for object detection and instance segmentation. *IEEE transactions on cybernetics*, 52(8): 8574–8586.
- 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*, 12846–12856.
- 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*, 5819–5828.