

SMamba: Sparse Mamba for Event-based Object Detection

Nan Yang^{1*}, Yang Wang^{1*}, Zhanwen Liu^{1†}, Meng Li², Yisheng An¹, Xiangmo Zhao¹

¹School of Information Engineering, Chang'an University, China

²School of Civil Engineering, Tsinghua University, China

{2022024001, ywang120, zwliu, aysm, xmzhao}@chd.edu.cn, mengli@tsinghua.edu.cn

Abstract

Transformer-based methods have achieved remarkable performance in event-based object detection, owing to the global modeling ability. However, they neglect the influence of non-event and noisy regions and process them uniformly, leading to high computational overhead. To mitigate computation cost, some researchers propose window attention based sparsification strategies to discard unimportant regions, which sacrifices the global modeling ability and results in suboptimal performance. To achieve better trade-off between accuracy and efficiency, we propose Sparse Mamba (SMamba), which performs adaptive sparsification to reduce computational effort while maintaining global modeling capability. Specifically, a Spatio-Temporal Continuity Assessment module is proposed to measure the information content of tokens and discard uninformative ones by leveraging the spatiotemporal distribution differences between activity and noise events. Based on the assessment results, an Information-Prioritized Local Scan strategy is designed to shorten the scan distance between high-information tokens, facilitating interactions among them in the spatial dimension. Furthermore, to extend the global interaction from 2D space to 3D representations, a Global Channel Interaction module is proposed to aggregate channel information from a global spatial perspective. Results on three datasets (Gen1, 1Mpx, and eTram) demonstrate that our model outperforms other methods in both performance and efficiency.

Introduction

Robust object detection is fundamental for intelligent systems, such as autonomous driving and robotics (Liu et al. 2024b; Zhang et al. 2013; Li et al. 2024; Liu et al. 2023). However, frame-based cameras encounter inherent limitations in frame rate and dynamic range, leading to low-quality images and hindering the extraction of discriminative features in challenging scenarios, such as high-speed motion and adverse exposure conditions (*e.g.*, low light and over-exposure) (Lichtsteiner, Posch, and Delbruck 2008; Sayed and Brostow 2021; Liu et al. 2024c; Yan, Gong, and Zhang 2018; Liu et al. 2024d). In these scenarios, event cameras

*These authors contributed equally.

†Corresponding author.

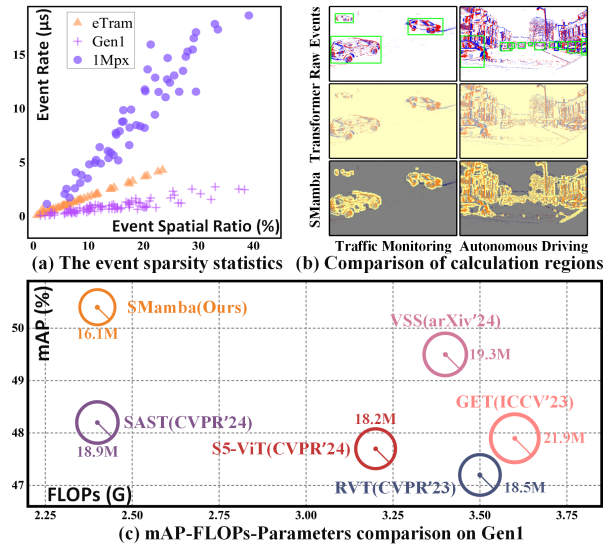


Figure 1: (a) The event sparsity statistics on three datasets. Each point represents a scenario. (b) Compared with calculating all regions uniformly in Transformer, our proposed SMamba can suppress non-event and noisy regions (gray regions) from participating in calculations and only retain information-rich areas (yellow regions), alleviating computational overhead significantly and suppressing noise disturbance, simultaneously. (c) mAP-FLOPs-Parameters comparison between state-of-the-art methods and our SMamba on Gen1, where the circle radius is the parameter. SMamba achieves a superior balance between accuracy and efficiency.

emerge as a superior alternative due to their unique advantages, which asynchronously detect pixel-level changes in light intensity, offering high temporal resolution and high dynamic range, exhibiting stable and robust performance under challenging conditions (Finatou et al. 2020; Huang et al. 2023; Son et al. 2017; Gallego et al. 2020).

To fully exploit the superior performance of event cameras in object detection task, researchers have developed various methods based on well-designed neural network architectures, including SNN-based, GNN-based, CNN-based, and Transformer-based methods. Theoretically, SNN-based (Cordone, Miramond, and Thierion 2022) and GNN-based

(Schaefer, Gehrig, and Scaramuzza 2022) methods can achieve low-latency inference, but require specialized hardware and perform poorly in real scenarios. Recent research indicates that Transformer-based (Gehrig and Scaramuzza 2023; Peng et al. 2023; Zubic, Gehrig, and Scaramuzza 2024; Zubić et al. 2023) methods, benefiting from their larger receptive field, significantly outperform local receptive field-based CNN methods (Dosovitskiy et al. 2021; Perot et al. 2020; Li et al. 2022). However, these methods typically transform event streams into discrete tokens and process non-event and noisy regions uniformly, neglecting the influence of spatial sparsity (Peng et al. 2024) and low signal-to-noise ratio (Ding et al. 2023), leading to substantial redundant computation and suboptimal performance. As shown in Figure 1(a), we calculate the ratio of triggered event pixels (event spatial ratio) and the average number of events over the camera’s operating time (event rate) for the Gen1 (De Tournemire et al. 2020), 1Mpx (Perot et al. 2020), and eTram (Verma et al. 2024) datasets. The results indicate that in 96% of the scenes, the event spatial ratio is below 30%, indicating high spatial sparsity. To improve the computation efficiency on sparse event data, SAST (Peng et al. 2024) proposes the window attention based sparsification strategy to discard unimportant regions but sacrifices the global modeling capability. Additionally, the token scoring module is optimized through an indirect gradient propagation path, complicating the optimization of the framework and necessitating more iterations.

In this paper, we propose Sparse Mamba (SMamba) which effectively reduces the computational cost by adaptively discarding uninformative tokens, and captures global context through information-guided spatial selective scanning and global spatial-based channel selective scanning, achieving better accuracy and efficiency trade-off, as depicted in Figure 1(b, c). Specifically, we design the Spatio-Temporal Continuity Assessment (STCA) module by leveraging the spatiotemporal distribution differences between activity and noisy events, without complex module designs and challenging optimizations. The information content of tokens is measured by evaluating the spatiotemporal continuity of events within each token’s corresponding region, generating a sparsification map to guide the non-event and noisy tokens discarding. Based on the assessment results, we propose the Information-Prioritized Local Scanning (IPL-Scan) strategy, which sorts tokens at the window level based on their information content, facilitating interactions among high-information tokens in 2D space while preserving critical local content. Additionally, to further extend the global interaction from 2D space to 3D representations, we propose the Global Channel Interaction (GCI) module to perform dynamic integration of channel information based on 2D global context. The experimental results on three datasets (Gen1, 1Mpx, and eTram) demonstrate that our method achieves superior performance and greater efficiency. Overall, our contribution can be summarized as follows:

(1) We propose the SMamba, which adaptively discards non-event and noisy tokens based on the spatiotemporal continuity assessment and captures the global relationship across spatial and channel dimensions, demonstrating an op-

timal balance between accuracy and efficiency.

(2) The IPL-Scan is designed to guide the model to focus on high-information tokens during the scanning process, thereby improving spatial contextual modeling capability.

(3) The GCI module is devised to extend global interactions into the 3D feature space by aggregating channel information from the global spatial perspective, further improving global modeling capability.

(4) Experimental results on Gen1, 1Mpx, and eTram datasets demonstrate that our SMamba surpasses state-of-the-art methods, achieving superior performance.

Related Work

This section provides a comprehensive overview of event-based object detection methods, followed by a review of recent advancements in Vision Mamba.

Event-Based Object Detection

Existing event-based object detection methods can be categorized according to the neural network architecture employed: SNN-based, GNN-based, CNN-based, and Transformer-based methods. SNN-based (Cordone, Miramond, and Thierion 2022) and GNN-based (Schaefer, Gehrig, and Scaramuzza 2022) methods effectively leverage the sparse and asynchronous properties of event data, thereby achieving high computational efficiency. However, the immaturity of their network architectures constrains their capacity to handle complex object detection task, leading to limited performance.

CNN-based and Transformer-based methods initially convert events into image-like tensors, facilitating compatibility with subsequent processing architectures. Prominent CNN-based methods, such as RED (Perot et al. 2020) and ASTMNet (Li et al. 2022), integrate CNNs with RNNs to effectively extract spatiotemporal features. Recently, Transformer-based methods (Gehrig and Scaramuzza 2023; Peng et al. 2023; Zubic, Gehrig, and Scaramuzza 2024) benefit from the global receptive field of self-attention mechanism, surpassing CNN-based methods and achieving outstanding performance. However, the self-attention mechanism uniformly processes non-event and noisy regions, resulting in considerable redundant computation. To address this, SAST (Peng et al. 2024) proposes a window-token co-sparsification strategy that adaptively discards insignificant windows and tokens. Nevertheless, this method relies on window attention to achieve low computational overhead, sacrificing global modeling capability.

Vision Mamba

Mamba achieves a better balance between efficiency and performance, making it an effective alternative to Transformer. Specifically, Mamba improves global modeling capability by introducing an input-dependent selective scanning mechanism (S6) and proposes a parallel scanning mechanism to retain the linear complexity of State Space Models (SSMs). Inspired by Mamba’s success in the NLP field, Vim (Zhu et al. 2024) and Vamba (Liu et al. 2024a) integrate Mamba into the design of vision backbone networks,

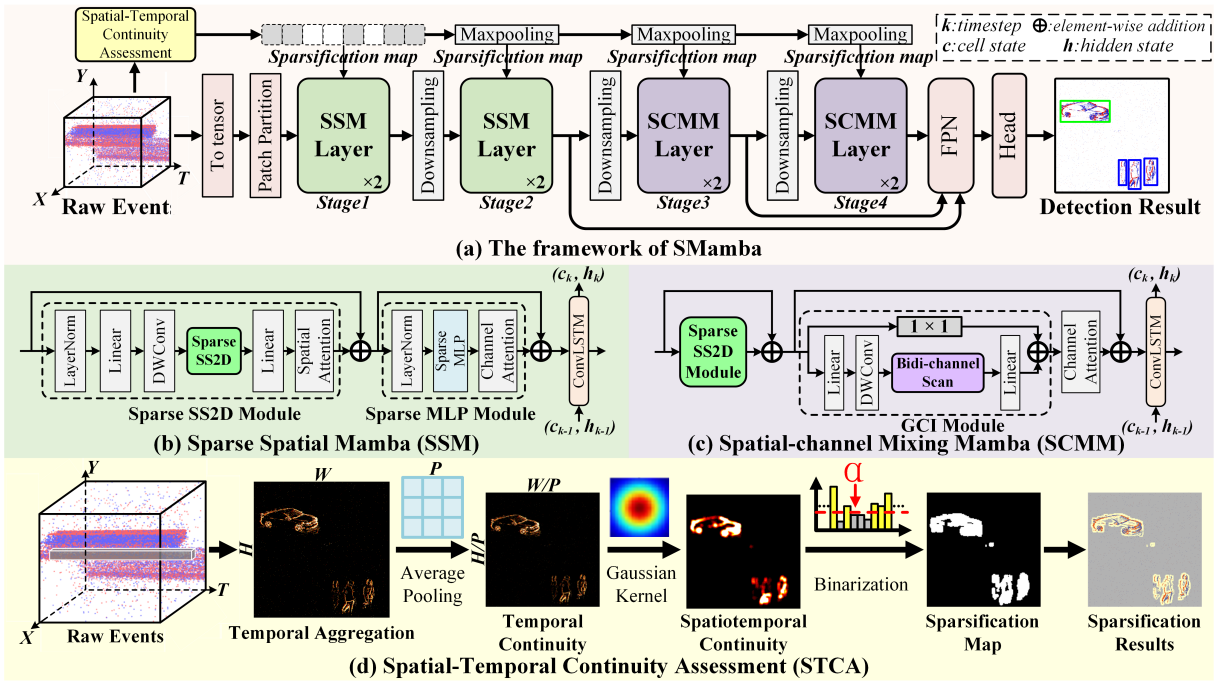


Figure 2: **The architecture of SMamba.** Given the input event stream, the STCA module (as illustrated in (d)) evaluates the information content of tokens based on spatiotemporal continuity and generates a sparsification map to guide the sparsification operation. Simultaneously, the event stream is preprocessed and transferred to four stages for comprehensive global interactions across spatial and channel dimensions. The first two stages employ the Sparse Spatial Mamba (SSM) layer (as detailed in (b)) to facilitate global spatial interactions on kept tokens. The final two stages utilize the Spatial-Channel Mixing Mamba (SCMM) layer (as depicted in (c)) to permute global modeling within the 3D representation space.

proposing multi-scan strategies to adapt to the non-causal nature of images, achieving groundbreaking results. Subsequently, numerous studies have applied Mamba to various vision tasks, including medical image segmentation (Xu et al. 2024), remote sensing image segmentation (Patro and Agneeswaran 2024), image restoration (Guo et al. 2024; Zou et al. 2024), and object detection and tracking (Huang et al. 2024; Dong et al. 2024), demonstrating Mamba’s significant potential in the visual domain. However, Mamba lacks effective mechanisms for handling sparse event data, resulting in substantial redundant computations.

Method

Preliminary: SSMs and Mamba

The classical State Space Model (SSM) is a linear time-invariant system. Given an input sequence $x(t) \in \mathbb{R}^L$, it maintains a hidden state $h(t) \in \mathbb{R}^N$ to store the contextual information and generates an output $y(t) \in \mathbb{R}^L$. The computational procedure is as follows:

$$h'(t) = \mathbf{A}h(t) + \mathbf{B}x(t), y(t) = \mathbf{C}h(t), \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the state matrix, $\mathbf{B} \in \mathbb{R}^{N \times L}$ and $\mathbf{C} \in \mathbb{R}^{L \times N}$ are the input and output matrix, respectively.

To improve the global modeling ability of SSM, Mamba (Gu and Dao 2023) proposes the input-dependent selective scanning mechanism (S6) which introduces nonlinearity and

discretization by incorporating a time scale parameter Δ and the zero-order hold (ZOH). The procedure is as follows:

$$h_t = \bar{\mathbf{A}}h_{t-1} + \bar{\mathbf{B}}x_t, y_t = \mathbf{C}h_t, \quad (2)$$

$$\bar{\mathbf{A}} = \exp(\Delta \mathbf{A}), \bar{\mathbf{B}} = (\Delta \mathbf{A})^{-1}(\exp(\Delta \mathbf{A}) - \mathbf{I}) \cdot \Delta \mathbf{B}. \quad (3)$$

The Overview of the SMamba

The framework of SMamba is illustrated in Figure 2(a). Initially, the event stream is input into the Spatial-Temporal Continuity Assessment (STCA) module, depicted in Figure 2(d), which generates a sparsification map to guide the sparsification operations. Meanwhile, the event stream is converted into voxel tensors (Zhu et al. 2019) and divided into patches for tokenization. These tokens are then processed through four stages for multi-scale feature extraction. The first two stages employ the Sparse Spatial Mamba (SSM) layer, as shown in Figure 2(b), which includes a Sparse SS2D (2D Selective Scan) module to improve global spatial interactions on kept tokens, a Sparse MLP module to further reduce computational overhead, and a ConvLSTM (Shi et al. 2015) to transmit spatiotemporal information across time steps, with its output sent to subsequent layers. The final two stages utilize the Spatial-Channel Mixing Mamba (SCMM) layer, as shown in Figure 2(c), which comprises a Sparse SS2D module, a Global Channel Interaction (GCI) module which extends the global modeling to 3D represen-




Figure 3: **Sparse SS2D**. For ease of observation, a large patch size is utilized for tokenization. The uninformative and noisy tokens (gray regions) are discarded from the calculation and the kept tokens (yellow regions) are expanded into three scanning sequences by Bidi-Scan and IPL-Scan, each sequence is processed in parallel using separate S6 blocks. The window sorting results of IPL-Scan are indicated at the center of each window (highlighted in red).

tations by facilitating channel interactions from a global perspective, and a ConvLSTM. Features produced by the last three stages are then fed into a Feature Pyramid Network (FPN) for multi-scale feature fusion. Finally, the YOLOX (Ge et al. 2021) detection head outputs the detection results.

Spatial-Temporal Continuity Assessment

Event cameras asynchronously trigger events at locations where brightness changes exceed the threshold, resulting in significant spatial sparsity, particularly in scenarios where the camera is stationary (Verma et al. 2024). Furthermore, the inherent circuitry characteristics of event cameras generate substantial noise (Ding et al. 2023; Duan 2024). These blank and noisy regions are uninformative, leading to unnecessary computations and potential interference.

We observe that activity events and noise events exhibit significant differences in their spatiotemporal distribution. Specifically, the noise events are spatially isolated or temporally discontinuous, whereas activity events are typically localized at the edges of moving objects, exhibiting spatial proximity and temporal continuity (Kim et al. 2021). Based on this prior, we propose the Spatial-Temporal Continuity Assessment (STCA) module, as depicted in Figure 2(d), which evaluates token importance and selectively discards uninformative tokens by assessing spatiotemporal continuity of events, reducing computational overhead.

Specifically, given the event stream $\{(x_i, y_i, t_i, p_i)\}_{i=1}^N$, where (x_i, y_i) is the spatial coordinates, t_i represents the timestamp, and $p_i \in \{-1, 1\}$ indicates the event polarity. The timestamps of events at each pixel are first accumulated to generate the temporal continuity score map $S^t \in \mathbb{R}^{H \times W}$, which quantifies the temporal continuity at each spatial location. The formulation is as follows:

$$S_{x,y}^t = \sum_{i, x_i=x, y_i=y} t_i. \quad (4)$$

Next, average pooling with kernel size and stride of P is employed to extract the temporal information content $S^t \in \mathbb{R}^{H/P \times W/P}$ corresponding to each token, where P represents the patch size utilized during event tokenization. Subsequently, the neighborhood information is effectively aggregated to assess spatial continuity. For activity events, nearby neighbors are more likely to be triggered by the same

moving edge, while distant neighbors are more likely to be noise. Therefore, to mitigate the impact of noise on information content evaluation, a Gaussian function is employed to perform distance-based weighted aggregation within the neighborhood, thereby smoothing the noise while maintaining more complete object structures (Wan et al. 2022). The formulation is as follows:

$$S^{st} = \frac{\sum_{q \in \Omega} \left(\exp\left(-\frac{\|q-c\|^2}{2\sigma^2}\right) S_q^t \right)}{\sum_{q \in \Omega} \exp\left(-\frac{\|q-c\|^2}{2\sigma^2}\right)}, \quad (5)$$

where c is the center of neighborhood Ω , S_q^t denotes the value of neighbor q , and σ represents the variance. In the resulting spatiotemporal continuity score map S^{st} , each pixel value indicates the activity event information content of the token. The larger the value, the more important the token.

The mean value of the spatiotemporal continuity score map S^{st} indicates the sparsity of the scene and serves as a threshold for discarding uninformative tokens. To adaptively retain important tokens based on the scene’s sparsity level and avoid the loss of critical object information, a manually adjusted sparsity scaling factor β is introduced to modulate the discard ratio. The threshold α is defined as follows:

$$\alpha = \frac{\text{sum}(S^{st})}{\beta \frac{HW}{P^2}}. \quad (6)$$

Based on this threshold, a sparsification map $D \in \mathbb{R}^{H/P \times W/P}$ is generated for sparsification. The expression is as follows:

$$D_{x,y} = \begin{cases} 1, & \text{if } S_{x,y}^{st} \geq \alpha, \\ 0, & \text{if } S_{x,y}^{st} < \alpha. \end{cases} \quad (7)$$

The sparsification map is subsequently propagated to the following layers to guide sparsification operations.

Information-Prioritized Local Scan

The 2D spatial scanning strategies, such as Bidi-Scan and Cross-Scan, may disperse tokens associated with the same object in the scanning sequence, resulting in distant scanning intervals and weakened interactions between them (Liu et al. 2024a; Shi, Dong, and Xu 2024). Therefore, we propose the

Information-Prioritized Local Scan (IPL-Scan) which mitigates the limitations of 2D scan methods, and design the Sparse SS2D which combines IPL-Scan and Bidi-Scan to facilitate global interactions, as depicted in Figure 3.

The spatiotemporal continuity score map quantifies the information of tokens, with higher scores indicating a higher likelihood of being a foreground object. Reordering tokens according to this map, tokens with higher informational content are processed earlier, which shortens the scan distance between important tokens, facilitating interactions among them. Furthermore, tokens with lower scores are processed later, which mitigates potential interference from noise.

Considering that the direct reordering might disrupt local information, a local constraint is introduced to the sorting process. When a token is processed, its $k \times k$ neighborhood is also processed immediately afterward. Specifically, max pooling with a kernel and stride of k is used to extract the maximum value from each $k \times k$ local window. These maximum values, representing the local windows, are initially sorted. Subsequently, the sorted results are upsampled by k to obtain the window-level sorted outcome. This strategy effectively facilitates the interaction among potential object regions while preserving local information.

Global Channel Interaction

To extend the global interaction from 2D space to 3D representations, we propose the Global Channel Interaction (GCI) module, as shown in Figure 2(c), which integrates the Bidirectional Channel (Bidi-channel) Scan mechanism with 1×1 convolution to perform dynamic integration of channel information based on global and local content.

The feature tensor $X \in \mathbb{R}^{C \times H \times W}$ is processed through two distinct branches: the Bidi-channel Scan and the 1×1 convolution, facilitating channel interactions at global and local levels. Within the Bidi-channel Scan branch, X is pre-processed by Linear and DWConv (Depthwise Convolution) to capture the local context and then sent into the Bidi-channel Scan, as shown in Figure 4. X is flattened along the H and W dimensions, followed by a transposition to treat the global spatial information of each channel as the fundamental unit of interaction. Subsequently, a reverse sequence is generated by flipping, which is then input into S6 along with the original sequence to enable adaptive interaction from the global perspective. Performing selective scanning based on global spatial content allows each channel to selectively focus on other channels from a more comprehensive perspective, accurately capturing dependencies between channels and further enhancing global modeling capability. The other branch employs 1×1 convolution to capture pixel-level dependencies between channels, enabling local adaptive interaction. Finally, the results from both branches are integrated to achieve comprehensive channel interactions.

Experiments

This section begins with an overview of the experimental setup. Subsequently, a comparative analysis of our method against state-of-the-art (SOTA) methods is presented. The visualization results are then shown to demonstrate the




Figure 4: **Bidirectional Channel Scan.** The feature map is unfolded along the spatial dimension and arranged into two sequences along the channel direction, which are then input into S6 for full spatial-level channel information interaction.

scene-adaptability of our method. Finally, ablation studies are conducted to validate the effectiveness of our approach.

Experimental Setup

This subsection details the employed datasets, the validation metrics utilized, and the implementation details.

Datasets. We conduct experiments on two autonomous driving datasets Gen1 (De Tournemire et al. 2020) and 1Mpx (Perot et al. 2020), and one traffic monitoring dataset eTram (Verma et al. 2024). The Gen1 dataset comprises over 39 hours of event data at 304×240 resolution, providing more than 255,000 labeled cars and pedestrians with annotation frequencies of 1 Hz, 2 Hz, or 4 Hz. The 1Mpx dataset provides 14.65 hours of data at a higher resolution of 1280×720 and 60 Hz annotated frequency, containing over 25 million labeled boxes across seven categories. The eTram dataset, collected for traffic monitoring, includes approximately 10 hours of data at 1280×720 resolution, encompassing around 2 million labeled boxes across eight categories with an annotation frequency of 30 Hz. The eTram is collected from a roadside perspective and exhibits greater sparsity compared to the other two datasets due to the camera’s fixed position (Verma et al. 2024).

Metrics. The COCO mAP (mean average precision) (Lin et al. 2014) is used to evaluate the accuracy of object detection. The model’s size is measured by the parameter count. Additionally, following SAST (Peng et al. 2024), we compute the average FLOPs (Floating Point Operations Per Second) on the first 1,000 samples of the test set to assess the computational complexity. We also compare the inference time (runtime) with other methods.

Implementation Details. To guarantee comparison fairness, we follow the dataset preprocessing methods, augmentation techniques, mixed batching strategy, event representation method and evaluation protocols established in RVT (Gehrig and Scaramuzza 2023).

Quantitative Results

We provide a comparative analysis of our method against 2 CNN-based methods: RED (Perot et al. 2020), ASTMNet (Li et al. 2022); as well as 5 Transformer-based methods: ERGO-12 (Zubić et al. 2023), RVT (Gehrig and Scaramuzza 2023), GET (Peng et al. 2023), SAST (Peng et al. 2024) and S5-ViT (Zubic, Gehrig, and Scaramuzza 2024) on the Gen1, 1Mpx datasets. On the eTram dataset, we compare

Methods	Backbone	Gen1				1Mpx			
		mAP	FLOPs	Params	Runtime	mAP	FLOPs	Params	Runtime
RED	CNN+RNN	40.0	6.0G	24.1M	16.7ms	43.0	19.0G	24.1M	39.3ms
ASTMNet	CNN+RNN	46.7	29.3G	>100M	35.6ms	48.3	75.7G	>100M	72.3ms
ERGO-12	Transformer	50.4	50.8G	59.6M	69.9ms	40.6	50.8G	59.6M	100ms
RVT-B	Transformer+RNN	47.2	3.5G	18.5M	10.2ms	47.4	10.3G	18.5M	11.9ms
GET-T	Transformer+RNN	47.9	3.6G	21.9M	16.8ms	48.4	10.6G	21.9M	21.9ms
SAST-CB	Transformer+RNN	48.2	2.4G	18.9M	22.7ms	48.7	6.4G	18.9M	23.6ms
S5-ViT-B	Transformer+SSM	47.7	>3.1G	18.2M	9.4ms	47.8	>9.1G	18.2M	10.9ms
VSS	SSM+RNN	49.5	3.4G	19.3M	17.4ms	48.2	10.3G	19.3M	17.7ms
Baseline	SSM+RNN	<u>50.0</u>	3.1G	16.1M	25.2ms	<u>48.8</u>	9.5G	16.7M	27.5ms
SMamba	SSM+RNN	50.4	2.4G (-23%)	16.1M	24.0ms	49.3	7.4G (-22%)	16.7M	26.0ms

Table 1: Performance compared with SOTA methods on two autonomous driving datasets Gen1 and 1Mpx. The reported FLOPs belongs to the backbone. Values in brackets (-) indicate the percentage decrease in FLOPs compared to the baseline method.

Methods	eTram			
	mAP	FLOPs	Params	Runtime
RVT-B	29.5	10.3G	18.5M	11.9ms
SAST-CB	30.0	6.2G	18.9M	24.4ms
S5-ViT-B	29.3	>9.1G	18.2M	10.9ms
VSS	31.3	10.3G	19.3M	17.7ms
Baseline	<u>32.3</u>	9.5G	16.7M	27.5ms
SMamba	32.6	6.6G (-31%)	16.7M	25.2ms

Table 2: Performance compared with SOTA methods on the traffic monitoring dataset eTram.

our method with 3 Transformer-based methods: RVT, SAST and S5-ViT, as the relevant code from other works has not been released. To compare with SSM-based methods, we employ the VSS block from VMamba (Liu et al. 2024a) to construct a detection framework named VSS. Moreover, a baseline model without the sparsification strategy is established to evaluate the effectiveness of the proposed method.

The results are shown in Table 1 and Table 2. On the Gen1 dataset, our SMamba outperforms all other methods with the lowest FLOPs and parameter count. Compared to ERGO-12, SMamba achieves the same mAP with merely 5% of its FLOPs and 27% of its parameter count. On the 1Mpx and eTram datasets, SMamba outperforms SAST-CB by 0.6% and 2.6% in mAP with similar FLOPs and lower parameter counts. By further integrating our sparsification strategy into our baseline, SMamba reduces FLOPs by 23%, 22% and 31%, along with mAP gains of 0.4%, 0.5% and 0.3% in three datasets, respectively. Our sparsification operation imposes network focus on important areas, alleviating interference from blank and noisy regions, thereby reducing computational overhead and improving accuracy. The inference speed of SMamba is faster than CNN-based methods and Transformer-based method ERGO-12, and is comparable with SAST-CB while achieving higher accuracy. The consistent performance improvements on the autonomous driving and traffic surveillance datasets demonstrate that our method can generalize to varying sparsity levels while achieving an ideal trade-off between accuracy and efficiency.

Methods	mAP	FLOPs	Params
Variance	30.8	6.5G	16.7M
Entropy	30.4	6.6G	16.7M
Scoring module	31.2	7.0G	17.0M
STCA	32.6	6.6G	16.7M

Table 3: Performance of different scoring methods.

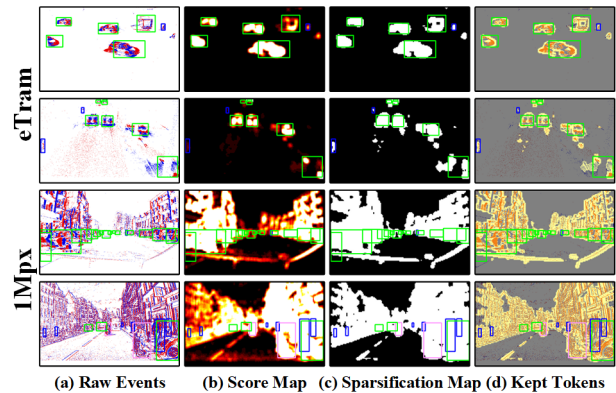


Figure 5: Visualizations of raw events, score map, sparsification map and sparsification results.

Sparsification Visualizations

The visualizations of raw events, score map, sparsification map, and sparsification results on the eTram and 1Mpx datasets are shown in Figure 5, with increasing levels of scene complexity. The eTram, collected by a stationary camera, exhibits greater sparsity compared to the 1Mpx which is obtained from a moving camera. As event density increases, the STCA module keeps an increasing number of tokens. This indicates that our STCA exhibits robust scene-adaptive capability, effectively mitigating interference from blank areas and noise while selecting important tokens.

Ablation Studies

To assess the effectiveness of the proposed method, we conduct a series of ablation studies on the eTram dataset.

Methods	mAP	FLOPs	Params
IPL-Scan	29.9	6.5G	16.3M
Bidi-Scan	30.2	6.6G	16.5M
Cross-Scan	30.7	6.6G	16.9M
Bidi-Scan + IPL-Scan	32.6	6.6G	16.7M
Bidi-Scan + IPL-Scan (w/o local constraint)	30.7	6.6G	16.7M

Table 4: Performance of different scanning patterns.

Bidi-channel Scan	1×1	mAP	FLOPs	Params
MLP		31.0	9.7G	19.2M
✓		31.5	6.3G	16.0M
✓	✓	32.6	6.6G	16.7M

Table 5: Performance of the components of the GCI module.

STCA Module. We compare our STCA with two information content evaluation metrics, variance and entropy, as well as a learnable event scoring module (Peng et al. 2024) while maintaining architectural consistency. As shown in Table 3, our method exhibits superior performance, surpassing other methods. Variance and entropy cannot distinguish between activity events and noise events, leading to misjudgments of information content. The scoring module is optimized through an indirect gradient propagation path, resulting in suboptimal token scoring. In contrast, the STCA effectively distinguishes activity events from noise and blank areas based on the spatiotemporal continuity prior.

Scanning Pattern. We perform an ablation on the scanning pattern within the Sparse SS2D, the results are shown in Table 4. Increasing the 2D scanning paths from 2 (Bidi-Scan) to 4 (Cross-Scan) improves the mAP by merely 0.5%. Employing IPL-Scan alone yields the poorest performance due to the significant loss of spatial structure information. However, the combination of Bidi-Scan and IPL-Scan achieves the best performance, increasing the mAP by 2.4%. Removing the local constraint in the IPL-Scan results in a significant decline in performance, indicating the crucial role of local information.

Figure 6 presents a visual comparison of our scan method with Bidi-Scan and Cross-Scan. Bidi-Scan and Cross-Scan organize scanning sequences in the 2D direction, which limits contextual interaction, resulting in less distinct features. In contrast, our IPL-Scan shortens the scanning distance between different regions of the same object, thereby facilitating interaction among these regions and enabling the modeling of more discriminative feature representations.

Global Channel Interaction Module Design. We replace the GCI module with MLP as the baseline to analyze the contribution of Bidi-channel Scan and 1×1 convolution. The results are shown in Table 5, replacing MLP with Bidi-channel Scan improves mAP by 0.5% while reducing FLOPs and parameters by 35% and 16.7%, respectively. The global spatial context provides a more comprehensive perspective for interaction, enabling more accurate channel information aggregation. Incorporating 1×1 convolution introduces pixel-level channel interactions, further




Figure 6: The feature visualization of Bidi-Scan, Cross-Scan and our method.

S1	S2	S3	S4	mAP	FLOPs	Params
				31.0	9.7G	19.2M
			✓	31.0	6.2G	16.7M
		✓	✓	32.6	6.6G	16.7M
	✓	✓	✓	31.4	6.8G	17.2M
✓	✓	✓	✓	31.5	7.0G	17.8M

Table 6: Performance of GCI placed after different stages.

increasing mAP by 1.1%, which suggests that local information is equally important for channel interactions.

Global Channel Interaction Module Placement. We use MLP as the baseline to examine the impact of placing the GCI module at different stages. The results in Table 6 suggest that placing the GCI module in the final two stages achieves the best performance, improving mAP by 1.6% while reducing FLOPs and parameters by 32% and 13%, respectively. This can be attributed to the higher semantic level and richer channel information in the last two stages, making comprehensive channel interactions more beneficial.

Conclusion

In this paper, we propose the Sparse Mamba which achieves a superior balance between accuracy and efficiency for event-based object detection. The STCA module discards non-event and noisy tokens adaptively, reducing computational overhead significantly. The IPL-Scan and GCI modules capture the global context in spatial and channel dimensions, respectively. The IPL-Scan shortens the scan distance among high-information tokens and facilitates spatial contextual modeling. The GCI module aggregates channel information from the global spatial perspective, achieving global interaction within 3D space. Experimental results on three datasets demonstrate that our method achieves superior performance and computational efficiency.

Limitation. Our STCA relies on the assumption of temporal and spatial independence of event noise, which may limit its effectiveness in handling large bursts or clusters of noise. We will further consider integrating more effective semantic learning to better distinguish spatiotemporally continuous noisy tokens in future work.

Acknowledgments

This work is supported in part by the National Key Research and Development Program of China under Grant 2023YFC3081700, in part by the National Natural Science Foundation of China under Grants U24B20127 and 52172302.

References

- 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.
- 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*.
- Ding, S.; Chen, J.; Wang, Y.; Kang, Y.; Song, W.; Cheng, J.; and Cao, Y. 2023. E-MLB: Multilevel benchmark for event-based camera denoising. *IEEE Transactions on Multimedia*, 26: 65–76.
- Dong, W.; Zhu, H.; Lin, S.; Luo, X.; Shen, Y.; Liu, X.; Zhang, J.; Guo, G.; and Zhang, B. 2024. Fusion-mamba for cross-modality object detection. *arXiv preprint arXiv:2404.09146*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Housley, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *ICLR*.
- Duan, Y. 2024. LED: A Large-scale Real-world Paired Dataset for Event Camera Denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25637–25647.
- Finatou, T.; Niwa, A.; Matolin, D.; Tsuchimoto, K.; Mascheroni, A.; Reynaud, E.; Mostafalu, P.; Brady, F.; Chotard, L.; LeGoff, F.; et al. 2020. 5.10 A 1280×720 back-illuminated stacked temporal contrast event-based vision sensor with 4.86 μm pixels, 1.066 GEPS readout, programmable event-rate controller and compressive data-formatting pipeline. In *2020 IEEE International Solid-State Circuits Conference-(ISSCC)*, 112–114. IEEE.
- 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.
- Ge, Z.; Liu, S.; Wang, F.; Li, Z.; and Sun, J. 2021. YOLOX: Exceeding YOLO Series in 2021. *arXiv preprint arXiv:2107.08430*.
- 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*, 13884–13893.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-Time Sequence Modeling with Selective State Spaces. *arXiv preprint arXiv:2312.00752*.
- Guo, H.; Li, J.; Dai, T.; Ouyang, Z.; Ren, X.; and Xia, S.-T. 2024. MambaIR: A Simple Baseline for Image Restoration with State-Space Model. In *ECCV*.
- Huang, J.; Wang, S.; Wang, S.; Wu, Z.; Wang, X.; and Jiang, B. 2024. Mamba-fetrack: Frame-event tracking via state space model. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, 3–18. Springer.
- Huang, K.; Zhang, S.; Zhang, J.; and Tao, D. 2023. Event-based simultaneous localization and mapping: A comprehensive survey. *arXiv preprint arXiv:2304.09793*.
- Kim, J.; Bae, J.; Park, G.; Zhang, D.; and Kim, Y. M. 2021. N-imagenet: Towards robust, fine-grained object recognition with event cameras. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2146–2156.
- Li, C.; Liu, Z.; Lin, S.; Wang, Y.; and Zhao, X. 2024. Intention-convolution and hybrid-attention network for vehicle trajectory prediction. *Expert Systems with Applications*, 236: 121412.
- Li, J.; Li, J.; Zhu, L.; Xiang, X.; Huang, T.; and Tian, Y. 2022. Asynchronous spatio-temporal memory network for continuous event-based object detection. *IEEE Transactions on Image Processing*, 31: 2975–2987.
- 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.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024a. VMamba: Visual State Space Model. *arXiv preprint arXiv:2401.10166*.
- Liu, Z.; Cheng, J.; Fan, J.; Lin, S.; Wang, Y.; and Zhao, X. 2023. Multi-modal fusion based on depth adaptive mechanism for 3D object detection. *IEEE Transactions on Multimedia*, 1–11.
- Liu, Z.; Li, C.; Wang, Y.; Yang, N.; Fan, X.; Ma, J.; and Zhao, X. 2024b. Multi-scale Temporal Fusion Transformer for Incomplete Vehicle Trajectory Prediction. *IEEE Transactions on Intelligent Vehicles*, 1–14.
- Liu, Z.; Li, Y.; Wang, Y.; Gao, B.; An, Y.; and Zhao, X. 2024c. Boosting Visual Recognition in Real-world Degradations via Unsupervised Feature Enhancement Module with Deep Channel Prior. *IEEE Transactions on Intelligent Vehicles*, 1–14.
- Liu, Z.; Yang, N.; Wang, Y.; Li, Y.; Zhao, X.; and Wang, F.-Y. 2024d. Enhancing traffic object detection in variable illumination with rgb-event fusion. *IEEE Transactions on Intelligent Transportation Systems*, 25: 20335–20350.
- Patro, B. N.; and Agneeswaran, V. S. 2024. Mamba-360: Survey of state space models as transformer alternative for long sequence modelling: Methods, applications, and challenges. *arXiv preprint arXiv:2404.16112*.

- 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*, 16794–16804.
- Peng, Y.; Zhang, Y.; Xiong, Z.; Sun, X.; and Wu, F. 2023. Get: Group event transformer for event-based vision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6038–6048.
- Perot, E.; de Tournemire, P.; Nitti, D.; Masci, J.; and Sironi, A. 2020. Learning to Detect Objects with a 1 Megapixel Event Camera. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 16639–16652. Curran Associates, Inc.
- Sayed, M.; and Brostow, G. 2021. Improved handling of motion blur in online object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1706–1716.
- 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*, 12371–12381.
- Shi, X.; Chen, Z.; Wang, H.; Yeung, D.-Y.; Wong, W.-K.; and Woo, W.-c. 2015. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems*, 28.
- Shi, Y.; Dong, M.; and Xu, C. 2024. Multi-Scale VMamba: Hierarchy in Hierarchy Visual State Space Model. *arXiv preprint arXiv:2405.14174*.
- Son, B.; Suh, Y.; Kim, S.; Jung, H.; Kim, J.-S.; Shin, C.; Park, K.; Lee, K.; Park, J.; Woo, J.; et al. 2017. 4.1 A 640×480 dynamic vision sensor with a 9μm pixel and 300Meps address-event representation. In *2017 IEEE International Solid-State Circuits Conference (ISSCC)*, 66–67. IEEE.
- Verma, A. A.; Chakravarthi, B.; Vaghela, A.; Wei, H.; and Yang, Y. 2024. eTraM: Event-based Traffic Monitoring Dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22637–22646.
- Wan, Z.; Wang, Y.; Tan, G.; Cao, Y.; and Zha, Z.-J. 2022. S2N: Suppression-Strengthen Network for Event-Based Recognition Under Variant Illuminations. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part III*, 716–733. Springer.
- Xu, R.; Yang, S.; Wang, Y.; Du, B.; and Chen, H. 2024. A survey on vision mamba: Models, applications and challenges. *arXiv preprint arXiv:2404.18861*.
- Yan, Q.; Gong, D.; and Zhang, Y. 2018. Two-stream convolutional networks for blind image quality assessment. *IEEE Transactions on Image Processing*, 28(5): 2200–2211.
- Zhang, Y.; Zhang, H.; Nasrabadi, N. M.; and Huang, T. S. 2013. Multi-metric learning for multi-sensor fusion based classification. *Information Fusion*, 14(4): 431–440.
- Zhu, A. Z.; Yuan, L.; Chaney, K.; and Daniilidis, K. 2019. Unsupervised event-based learning of optical flow, depth, and egomotion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 989–997.
- Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*.
- Zou, W.; Gao, H.; Yang, W.; and Liu, T. 2024. WaveMamba: Wavelet State Space Model for Ultra-High-Definition Low-Light Image Enhancement. In *ACM Multimedia 2024*.
- 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 (ICCV)*, 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.