Compact SSM Hidden States for Accurate and Efficient Event-Based Vision

Nikola Zubić[1] Haiqing Hao[2] Davide Scaramuzza[1]

[1] Robotics and Perception Group, University of Zurich

[2] Department of Precision Instrument, Tsinghua University

Abstract

Event cameras record brightness changes asynchronously, as a sparse stream of events. They offer high temporal resolution and low power consumption, which makes them a natural sensor for detecting moving objects. But accurately detecting objects in sparse events data is compute expensive. To address this problem, many methods first group the events from a short time interval into packets, i.e., image-like tensors that show where and when activity occurred. Processing packets densely keeps rich spatial context and thus gives high accuracy, but exchanging information across all spatial locations at every layer makes the energy cost high. Event-by-event methods instead update only a small part of the network per event. This keeps the energy cost low, but sacrifices accuracy. We combine the best of both worlds, achieving high accuracy and low energy cost. Our detector processes event packets, but first compresses their many spatial tokens into a few compact state-space model (SSM) hidden states. The main feature mixing runs only on these states, and the result is expanded back to every location to find objects. On three public benchmarks (Gen1, N-Caltech101 and DSEC), our detector improves accuracy over the strongest asynchronous baselines from 0.72 to 9.92 AP points, while using 20× less computation and estimated energy. This shows that compact SSM hidden states make accurate event-based detection energy-efficient.

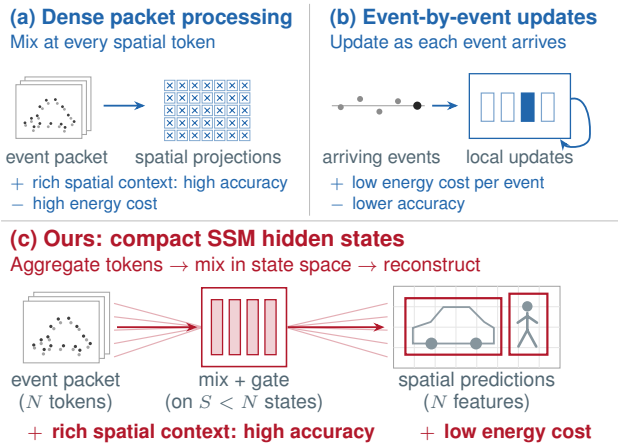


Figure 1. **Best of both worlds: rich spatial context at low energy cost. (a)** Dense packet methods collect events into an image-like tensor and apply channel mixing at every spatial location. This keeps rich spatial context and thus high accuracy, but repeating the mixing at every location makes the energy cost high. **(b)** Event-by-event methods update only a small part of the network when each event arrives. This keeps the energy cost per event low, but accuracy is lower. **(c)** Our detector processes event packets, but it first compresses the many spatial tokens of a packet (N) into a much smaller set of compact SSM hidden states (S). The mixer's main channel mixing and gating run only on these few states. The result is then expanded back to all N locations, so the detector can still locate objects. We therefore get the best of both worlds: the rich spatial context of packets, and thus high accuracy, at a low energy cost.

1. Introduction

Standard cameras, like those used in automotive, record a new image only every 33 ms. Between two frames, a car driving at 50 km/h moves about 0.4 m, and a pedestrian can step onto the road, unseen until the next frame arrives [11]. Event cameras close this gap. Instead of recording frames, each pixel independently reports a change in brightness the moment it occurs as an event, with a position, a timestamp, and a polarity (the sign of the change). Detecting objects with an event camera therefore means turning this asynchronous stream of events into bounding boxes that say where each object is and what it is.

This sensing principle gives event cameras submillisecond latency, microsecond temporal resolution, high dynamic range, and low power consumption. They see fast motion without blur and keep working when a scene is too dark or too bright for standard cameras [11]. These strengths, however, only matter if the detector can keep up. A detector that is accurate but expensive cannot run often enough on a robot or vehicle with a limited energy budget. A detector that is cheap but inaccurate misses the objects that matter. Event-based detection is only useful when it is both accurate and efficient.

This problem is difficult because a single event carries little information. Objects only become visible when many events are combined across space. Gathering this spatial context makes a detector accurate, but also expensive.

Existing work follows two directions, shown in Fig. 1. Packet-based methods group the events from a short time window into a packet: an image-like tensor that records where and when activity occurred [9, 14, 42]. A standard network then processes this tensor densely. This keeps rich spatial context and gives high accuracy, but the network mixes features at every spatial location, so the energy cost is high. Event-by-event methods instead update only the part of the network affected by each new event, using graphs [11, 48] or sparse linear attention [18]. This keeps the energy cost per event low, but their accuracy remains below that of leading packet-based models. Moreover, when an application needs a result only once per packet, event-by-event processing cannot recover the latency already spent waiting for the packet, so processing the packet as a whole is the natural choice. The next frontier is therefore to process packets not only accurately, but also at the amortized energy cost of event-by-event updates.

The stumbling block is channel mixing. In a packet-based network, much of the work consists of projections and gates applied to the feature vector at every spatial location, so this work grows with the size of the packet. Faster spatial mixing, such as linear-time attention or SSMs, does not remove it, because every location is still transformed. One could avoid it by pooling the whole packet into a single vector, as image classifiers do with global average pooling. But such a vector only tells us what is in the scene, not where it is, so it cannot be used to place bounding boxes with high accuracy.

We present a family of packet-based event detectors built on compact state-space model (SSM) hidden states. We start from the hidden-state mixer of EfficientViM [26], which builds on Mamba-2's state-space duality [7], and adapt it to event packets. We combine it with event statistics, multiscale prediction and task-specific training. For RGB–event detection on DSEC [15], we also introduce conditional inference. The full detector runs when a new RGB frame arrives and saves its temporal state. A cheaper branch answers queries between frames by reusing the RGB features and the saved state, without changing it. Each query thus uses the newest events without running the full detector.

The key idea is to first compress the information, then mix it. Each mixer gathers the N spatial tokens of a packet into a few hidden states ($S < N$), runs its main channel mixing and gating on these states, and then expands the result back to all N locations. These transformations therefore act on S vectors instead of N, while the detector still sees a full spatial map to locate objects. The packet keeps its rich spatial context, but the cost of mixing no longer grows with the number of locations.

On Gen1, our best detector reaches 39.65 AP, which is 2.15 points above SSLA-L and 7.55 points above DAGR-L. All four model sizes exceed SSLA-L. On N-Caltech101 Detection, it reaches 61.42 AP, 9.92 points above SSLA-L, while using 17.15× less computation and 18.52× less estimated energy [18]. On DSEC, our RGB–event detector improves endpoint and interframe AP over DAGR by 1.13 and 0.72 points. It needs 0.56 MFLOPs and 0.45 μJ per event, which is 11.48× and 12.02× less than DAGR's published figures, and its computation is also below the 0.57 MFLOPs/event of DAGR's cheaper directed variant [11].[1] The same backbone also classifies events, with 91.91% and 69.68% top-1 accuracy on N-Caltech101 and N-ImageNet. On Gen1, pretraining adds 7.02 AP over training from scratch, and polarity adds 1.58 AP.

Our contributions are:

- A family of packet-based event detectors that compress spatial tokens into compact SSM hidden states before channel mixing. This gives high accuracy at a low computation and energy cost.
- Conditional SSM inference for RGB–event detection, which keeps full detections when RGB frames arrive and answers frequent queries between them with a cheaper branch.
- Four model sizes evaluated on three public detection benchmarks, supported by classification results and pretraining and polarity ablations.

We report DSEC semantic segmentation experiments in the supplementary material.

2. Related Work

Event representations and packet-based perception. Aggregating events into spatial tensors makes their irregular measurements compatible with established visual recognition architectures. Event Spike Tensors provide a differentiable framework for learning this aggregation jointly with a downstream task [12]. ERGO-12 selects informative representation channels using a Gromov–Wasserstein discrepancy between raw events and their encoded features [69]. EventPillars encodes temporal range, polarity and normalized event density in a compact representation, reducing encoding costs while retaining useful scene structure [9]. For detection, RED combines dense event representations with convolutional features and recurrent temporal memory [42]. RVT extends this approach with local and global spatial attention and recurrent feature aggregation [14]. We retain packet statistics and multiscale features, and reduce the cost

[1] Ratios compare our amortized costs with published reference figures. DSEC accounting covers complete streams at about 192 queries/s. N-Caltech101 uses one prediction per clip. Energy is estimated at 1.69 pJ/MAC.

after encoding by moving the mixer's main projections and gating into compact SSM hidden states.

Asynchronous and sparse inference. A complementary direction reduces the work triggered by each arriving event. AEGNN maintains a spatiotemporal graph and updates activations only at nodes affected by new measurements [48]. FARSE-CNN combines recurrent processing with sparse convolutions and hierarchical compression to process asynchronous event streams [47]. DAGR combines an image CNN with an asynchronous event GNN, injecting image context into the event branch to support predictions between image arrivals [11]. SSLA partitions linear-attention state into spatially selective components and combines sparse recurrent updates with parallel scatter–compute–gather training [18]. Spiking detectors, including MSD, pursue efficiency through spiking feature extraction and multiscale detection [31]. These approaches exploit different forms of sparsity and state reuse. We study an additional route: lowering the cost of processing event packets through compact-state computation. Our DSEC detector also reuses RGB features and committed temporal state, while retaining current-event evidence at each scheduled query.

State-space models for visual and event processing. Mamba introduces input-dependent selection in state-space sequence models, with computation that scales linearly with sequence length [17, 50, 68, 71]. Mamba-2 develops state-space duality, connecting structured SSMs and attention and enabling efficient mixing algorithms [7, 67]. Vision Mamba adapts selective SSMs to images through bidirectional sequence processing [43, 66]. Earlier event-camera models combine temporal SSMs with spatial attention and learnable timescales to improve robustness to changes in inference frequency [70]. SMamba exploits continuity for token sparsification, combines spatial and channel interactions, and uses ConvLSTM modules for temporal memory [57]. FocusMamba uses motion information for collaborative RGB–event token sparsification and fusion [58]. Building on compact hidden-state mixing [26], we adapt it to packet-based event detection. Spatial tokens are aggregated before the mixer's main channel projections and gating, then reconstructed into features for localization. This retains spatial evidence from event packets while enabling accurate detection at low amortized computation and MAC-energy cost.

Multimodal alignment, routing, and supervision. RGB–event fusion must reconcile fresh event measurements with appearance features from an earlier frame. FAOD learns event-guided alignment from temporally shifted training pairs [63]. ACGR aligns frame and event graphs into a shared representation [27]. HDI-Former couples an image ANN and an event SNN through dynamic cross-modal interaction [28]. Adaptive fusion and temporally refined pseudo-labels further support detection across operational frequencies [35]. These approaches motivate treating multimodal alignment, query routing, and supervision as distinct design choices. Our conditional SSM detector routes acquisition-boundary queries through the full prediction path and interior queries through a cheaper branch. Each interior query uses causal event prefixes and previews committed memory. State is committed once per acquisition. This separates query frequency from memory updates and reduces repeated inference work.

Pretraining and transfer for event vision. Limited event supervision makes initialization and representation adaptation important [36]. N-ImageNet demonstrates the benefits of large-scale event pretraining for recognition and learning with few labeled examples [23]. Spatiotemporal prompting transfers image-pretrained knowledge by adapting sparse event representations while preserving their spatial and temporal structure [32]. Our event-vision family combines pretrained SSM features with task-specific event statistics and detection heads. The pretraining and polarity ablations examine these contributions separately from the choice of mixer. Together with the classification evaluations, they test how compact SSM features transfer across event tasks, while detection establishes the central accuracy–computation and MAC-energy comparison.

3. Method

Our architecture retains the spatial context of event packets while reducing the cost of global mixing. Figure 2 shows its two uses: event-only recognition and detection, and causal RGB–event detection with frequent queries. Both use compact SSM hidden states. The latter additionally reuses image features and temporal memory across queries.

3.1. From event packets to compact hidden states

An event $e_i = (x_i, y_i, t_i, p_i)$ records a pixel's brightness change at time t_i with polarity p_i. We group events into five temporal bins. Each stores the first and last event times, polarity balance, and normalized event counts [9]. These 20 channels describe when, where, and how much activity occurred. Empty cells are zero. Gen1 and N-Caltech101 Detection bin at most 50,000 recent events over their time span. DSEC uses fixed 10-ms bins starting at the previous RGB timestamp, allowing successive queries to reuse accumulated statistics.

A convolutional stem feeds three spatial SSM stages. For event-only detection, S, B, and M use two blocks per stage (six total), while L uses $(3, 4, 2)$ blocks (nine total). Each block contains one mixer with its own parameters. Every mixer gathers spatial tokens into compact states, transforms

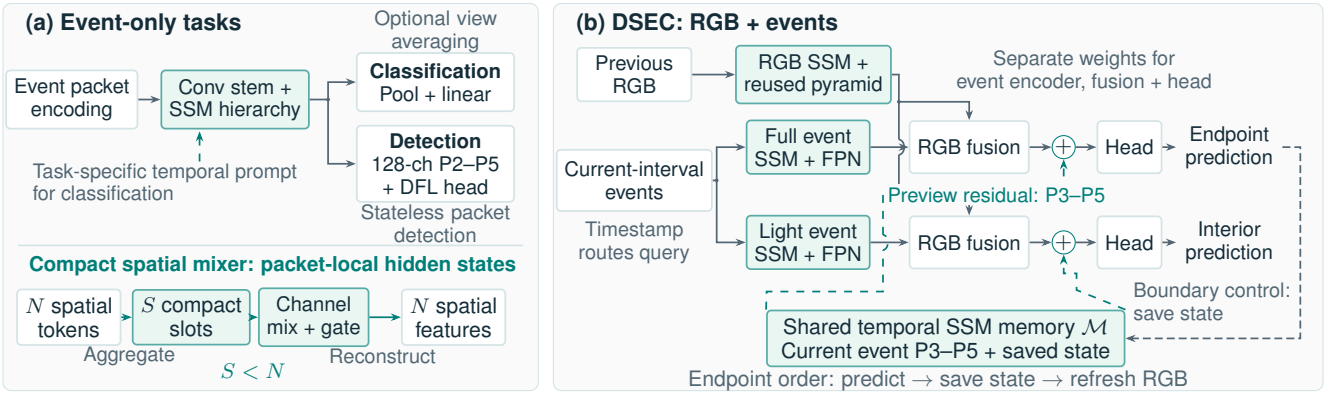


Figure 2. **Compact SSM architectures for event vision. (a)** A convolutional stem and SSM stages turn encoded event packets into spatial features. Pooling supports classification, while a multiscale head predicts object boxes for Gen1 and N-Caltech101 Detection. The lower diagram shows tokens gathered into compact hidden states for channel mixing, then reconstructed as spatial features. **(b)** DSEC combines current events with reused features from the previous RGB frame. The full path predicts at frame boundaries, saves temporal state, then refreshes RGB features. The lighter path predicts between boundaries without changing the saved state.

them, and reconstructs spatial features.

Within each mixer, $X \in \mathbb{R}^{C\times N}$ contains N spatial tokens with C channels. A pointwise projection and depthwise convolution predict write, read, and normalization maps $B, R, \Delta \in \mathbb{R}^{S\times N}$. Following compact state-space mixing [7, 26], we compute

$$\begin{aligned} A &= \mathrm{softmax}_N(\Delta), \quad H = X(A \odot B)^\top, \\ (U, Z) &= P_h(H), \\ Y &= P_o(U \odot \mathrm{SiLU}(Z) + dU)\, R. \end{aligned} \tag{1}$$

Here H contains S hidden states. The projections P_h, P_o and gate transform their channels, the learned scalar d scales a residual contribution, and R reconstructs spatial features for localization.

The saving comes from transforming fewer vectors. At fixed expansion ratio, a channel projection costs $O(C^2)$ per vector. Applying it to S hidden states instead of N spatial tokens changes its cost from $O(NC^2)$ to $O(SC^2)$, reducing computation when $S < N$. Gathering and reconstructing spatial features costs $O(NCS)$. Depthwise convolutions and feed-forward layers still operate spatially. These hidden states are recomputed for each packet, distinct from the persistent temporal memory below.

3.2. Event-only prediction

For N-Caltech101 classification and event-only detection, EventPillars packet encoding in Figure 2(a) converts temporal, polarity, and density statistics into a spatial tensor. The convolutional stem extracts local patterns, and the SSM hierarchy mixes information across locations. The lower diagram illustrates its compact-state mixer (Sec. 3.1). Task-specific heads produce predictions.

For classification, pooling averages each stage's hidden states and the final spatial features into summary vectors. Linear classifiers turn these vectors into class scores, which we combine. N-Caltech101 uses one encoded recording. N-ImageNet instead uses three overlapping 30-ms views, each encoded as five bins of positive and negative event counts. Its temporal SSM prompt is an input adapter, a trainable module that converts event bins into features the pretrained backbone can use. It processes the bins in time order, using a hidden state to retain earlier information. Residual SSM adapters then add learned corrections to the backbone features. The same network processes all views, and averaging their class scores produces one prediction.

For detection, spatial locations are retained to predict object positions. A feature pyramid combines fine and coarse features into four maps, called P2, P3, P4, and P5. Their height and width are reduced from the padded input by factors of 4, 8, 16, and 32, respectively. Each has 128 channels. Fine maps preserve small-object detail, while coarse maps provide broader context. A shared head predicts class scores. Its box branch learns four edge-distance distributions using distribution focal loss (DFL). Their probability-weighted means define box edges relative to grid-cell centers. Non-maximum suppression removes overlapping duplicates. As noted above, Gen1 and N-Caltech101 Detection process packets independently, carrying no detector memory between them. This is the stateless path in the figure. The separate DSEC event-only experiment adds temporal SSM processing of the current packet and two preceding packets, as described in Sec. 4.2.

3.3. Causal RGB–event detection

Figure 2(b) adds RGB appearance and temporal memory. Let RGB frames arrive at t_k. A query $q \in (t_k, t_{k+1}]$ uses the previous image I_k and events with $t_k \le t_i < q$. The RGB SSM computes its pyramid once per frame, then reuses these

features between arrivals. At a frame-arrival boundary, the full event path predicts, saves memory, and refreshes RGB features for the next interval. Queries between boundaries use a lighter private path with its own event encoder, fusion, and head. They form temporary memory updates without saving them, so additional queries cannot change subsequent predictions. Empty and unlabelled intervals also execute boundary updates.

At each pyramid level, a learned gate combines shared RGB–event features with features specific to each modality. RGB age and event activity guide their relative weights. The RGB-only terms are reused, while event-dependent terms are recomputed for each query.

Temporal memory consists of two SSM banks at P3–P5. Each bank uses 32, 48, and 64 SSM state channels per spatial location at these respective levels. Every channel has four memory components with learned decay times, initialized to 10, 50, 250, and 1,000 ms. Short decay times allow quick responses to new events, while long ones retain older context. From the last saved state h_k, each component computes

$$\widetilde{h}_q = e^{-(q-t_k)/\tau} \odot h_k + \left(1 - e^{-(q-t_k)/\tau}\right) \odot u(E_q), \quad (2)$$

where τ is its decay time and $u(E_q)$ transforms the current event features into new evidence. Learned readouts turn both banks' states into corrections added to the fused features. Only the full boundary path saves $h_{k+1} = \widetilde{h}_{t_{k+1}}$.

Both heads predict distances from each P2–P5 grid point to the four box edges. Each distance uses 17 bins numbered 0–16, hence $\texttt{reg_max} = 16$. Adjacent bins are one pyramid stride apart, for example eight pixels at P3. The head averages these candidate distances using their predicted probabilities as weights.

The lighter path keeps 128-channel outputs but uses fewer blocks and narrower intermediate layers. Compressed projections mix features through fewer channels. When input and output widths match, a scaled copy of each input channel is also added. Spatial compression uses three linear convolutions: 1×1 reduces channels, 3×3 filters neighbouring locations at this smaller width, and 1×1 restores 128 channels. The expensive spatial filtering therefore operates on fewer channels. A parallel 3×3 filter processes each channel separately, and the two outputs are added. Only this lighter path varies across DSEC's four sizes. The full endpoint detector and both memory banks stay fixed.

3.4. Learning and amortized execution

Varifocal loss [62] learns class scores reflecting box quality, GIoU [46] improves box overlap, and DFL [29] learns edge-distance distributions. Their weights are 1, 2.5, and 0.5.

For DSEC, the frozen full detector teaches the lighter branch using the same available events, RGB frame, and saved history. Cross-head distillation [53] feeds student features after the first head block through the teacher's remaining layers. Matching their predictions and internal features guides learning without adding inference work. Boxes are imitated only when the teacher localizes better against ground truth.

Between annotated frames, we interpolate boxes along matching object tracks and supervise only their known classes. Unknown locations and classes receive no negative labels. Fully labelled endpoints also teach the lighter branch to reject background.

We count all DSEC inference work: RGB refreshes, encoding, every query, memory updates, and postprocessing. With M multiply–accumulates (MACs), A scalar operations, and N_e unique filtered events, amortized cost is $(2M + A)/(10^6 N_e)$ MFLOPs/event. Each MAC counts as two FLOPs. At 1.69 pJ/MAC, the energy proxy is $1.69M/(10^6 N_e)$ μJ/event. These averages charge all repeated work while counting each input event once.

4. Experiments

We test whether packet-based SSMs can improve detection accuracy while keeping inference inexpensive. We first compare four model sizes with asynchronous and dense detectors, then examine their computational costs and the training choices that contribute to accuracy. We also evaluate event classification to test whether the same compact SSM design can recognize object categories from event packets.

4.1. Experimental setup

Datasets and metrics. Gen1 [8] uses events alone to detect cars and pedestrians on the official test split, with a warm-up of 0.5 s and the standard small-box filtering. N-Caltech101 Detection [39] uses 906 held-out clips and 100 foreground classes, with one prediction at the end of each clip. DSEC uses causal events, either alone or with the previous RGB frame. RGB+event evaluation uses six validation recordings and two classes, car and pedestrian. Endpoint AP covers 5,897 labelled frame boundaries, while interframe AP covers 5,094 released intermediate targets. The two scores use different annotated samples. Their difference therefore does not directly measure whether detection becomes more or less accurate between RGB frames. Detection AP averages IoU thresholds from 0.50 to 0.95. We report AP50 separately and express all accuracies in percent.

Models and selection. For event-only detection, including DSEC, S/B/M/L denote four spatial SSM sizes. Only RGB+event DSEC uses four lighter branches sharing one frozen endpoint detector. Training follows Sec. 3.4, with pretrained SSM backbones, AdamW [34], and validation-based selection. Published baselines retain their reported evaluation protocols, identified in the tables.

Table 1. **Event-only detection on Gen1 and N-Caltech101.** All four compact SSMs exceed SSLA-L in AP on both datasets. AP and AP_{50} are percentages. F denotes computation in MFLOPs/event and E denotes MAC energy in μJ/event. Bold and underlining mark the best and second-best listed values in each column, respectively, including ties.

	Gen1 Detection				N-Caltech101 Detection			
Method	AP ↑	AP_{50} ↑	F ↓	E ↓	AP ↑	AP_{50} ↑	F ↓	E ↓
Dense detectors (P)								
Inception+SSD [20]	30.10	–	27,183	23,502	–	–	–	–
Events+RRC [4]	30.70	–	21,758*	18,386	–	–	–	–
MatrixLSTM+YOLOv3 [1]	31.00	–	34,519*	29,168	–	–	–	–
Events+YOLOv3 [21]	31.20	–	65,558	55,397	–	–	–	–
RED [42]	40.00	–	4,712	3,982	–	–	–	–
S5-ViT-B [70]	47.70	–	–	–	–	–	–	–
EVA+RVT-B [19]	47.70	–	3,500	–	–	–	–	–
GET [40]	47.90	–	3,600	–	–	–	–	–
SAST-CB [41]	48.20	–	2,400	–	–	–	–	–
SMamba [57]	<u>50.40</u>	–	2,400	–	–	–	–	–
ERGO-12 [69]	<u>50.40</u>	–	50,800	–	–	–	–	–
EventPillars [9]	**53.10**	–	50,800	–	–	–	–	–
Asynchronous detectors (P)								
NVS-S [30]	–	8.60	7.80	6.59	–	34.60	7.80	6.59
AsyNet [37]	–	14.50	205.00	173.00	–	64.30	200.00	169.00
AEGNN [48]	–	16.30	5.26	4.44	–	59.50	7.41	6.26
FARSE-CNN [47]	–	30.00	0.14	–	–	–	–	–
EHGCN [3]	–	–	–	–	–	69.40	1.06	–
DAGR-N [11]	26.30	–	1.36	1.15	–	62.90	2.28	1.93
DAGR-S [11]	30.40	–	4.58	3.87	–	70.20	6.85	5.76
DAGR-M [11]	31.80	–	9.94	8.40	–	72.70	12.20	10.60
DAGR-L [11]	32.10	–	17.40	14.70	–	73.20	18.90	16.00
SSLA-S [18]	33.40	62.90	0.10	0.09†	44.40	68.10	0.13	0.11†
SSLA-B [18]	35.10	65.50	0.18	0.15†	48.30	72.00	0.23	0.20†
SSLA-M [18]	37.00	67.00	0.41	0.34†	49.50	72.40	0.52	0.44†
SSLA-L [18]	37.50	67.50	0.72	0.61†	51.50	74.30	0.93	0.78†
Our compact SSMs (Gen1 W / N-Caltech101 T)								
Ours S	38.24	67.24	**0.05**	**0.04**	55.98	77.97	**0.04**	**0.03**
Ours B	37.73	66.84	<u>0.06</u>	<u>0.05</u>	58.45	80.61	<u>0.04</u>	<u>0.03</u>
Ours M	39.65	**69.39**	0.08	0.07	**61.42**	**83.79**	0.05	0.04
Ours L	39.04	<u>67.66</u>	0.09	0.08	<u>61.32</u>	<u>82.93</u>	0.06	0.05

Cost definitions. P retains published native costs from SSLA [18] and DAGR [11]. W counts supported neural-forward operations per historical packet event, excluding representation construction and unsupported operations. T counts the complete terminal prediction per unique filtered event. * marks a backbone-only lower bound. † marks converted energy, $E = 0.845F$, using two FLOPs/MAC and 1.69 pJ/MAC. Dashes indicate unreported results.

4.2. Object detection

Gen1 automotive detection. Table 1 shows that every model size exceeds the published 37.50 AP of SSLA-L. Our medium model reaches 39.65 AP and 69.39 AP50, improving AP by 2.15 points over SSLA-L and 7.55 points over DAGR-L. Even S reaches 38.24 AP. These gains show that compact-state processing preserves useful localization information at modest model sizes. Increasing capacity is not uniformly beneficial: M outperforms L, so we report the complete family rather than only its largest member.

The same table includes dense approaches from the published comparisons. Our models exceed Inception+SSD, Events+RRC, MatrixLSTM+YOLOv3, and Events+YOLOv3 in AP. Some dense systems achieve higher accuracy, including EventPillars at 53.10 AP, but their reported computational costs are substantially higher.

N-Caltech101 Detection. The improvement also transfers beyond automotive scenes. Our M model reaches 61.42 AP and 83.79 AP50, exceeding SSLA-L by 9.92 AP and 9.49 AP50 points. Its AP50 also exceeds DAGR-L by 10.59 points. Across corresponding S/B/M/L sizes, AP gains over SSLA are 11.58, 10.15, 11.92, and 9.82 points. The two detector seeds give mean AP and standard deviation of 55.93 ± 0.08, 57.67 ± 1.10, 61.38 ± 0.05, and 61.25 ± 0.10, respectively.

DSEC detection between RGB frames. We compare packet-based SSMs with asynchronous DAGR on the same validation targets. L reaches 43.41 endpoint AP and 47.01 interframe AP, exceeding DAGR-R50 by 1.13 and 0.72 points (Table 2(a)). All four RGB+event variants use the same frozen endpoint detector and weights, so their endpoint predictions and AP are identical. S/B/M/L vary only the cheaper branch used between RGB frames. Interframe AP rises from 35.85 (S) to 47.01 (L), showing the benefit of retaining capacity for current-event processing between RGB arrivals.

The added references retain their two-class test protocols. EFH-18 (†) reports AP at 25 ms, while EFH-10, Events+YOLOX and DAGR-R18/R34 (‡) report peak AP over 0–50 ms [11, 61]. SpikeYOLO8l uses CNN fusion and its native late-interframe evaluation (§) [22]. Our models

Table 2. **DSEC Detection.** (a) Two-class RGB+event detection. (b) Our held-out test TTA results and reported event-only baselines [18]. TTA combines flips and, for S, multiple scales, selected on validation. AP/AP$_{50}$ are percentages. Costs apply to (a): F is MFLOPs/event and E is arithmetic μJ/event. Markers are defined below. Bold/underlining mark best/second-best; AP rankings in (a) cover common validation only.

(a) RGB+event detection

Method	Endpoint AP ↑	Interframe AP ↑	F ↓	E ↓
EFH-10 [61]	–	41.90‡	1.88*	1.79*
EFH-18 [61]	–	45.10†	3.32*	3.09*
SpikeYOLO8l [22]	–	30.70§	–	–
Events+YOLOX [11]	40.20	42.50‡	22,049^{P}	18,631^{P}
DAGR-R18 [11]	37.60	39.70‡	6.74^{P}	5.70^{P}
DAGR-R34 [11]	39.00	41.40‡	6.57^{P}	5.55^{P}
DAGR-R50 [11]	<u>42.28</u>	<u>46.29</u>	6.41^{P}	5.42^{P}
Ours S	**43.41**	35.85	**0.39**	**0.31**
Ours B	**43.41**	42.93	<u>0.44</u>	<u>0.35</u>
Ours M	**43.41**	45.54	0.49	0.39
Ours L	**43.41**	**47.01**	0.56	0.45

DAGR-R50 and ours: common validation. Other rows: native test AP, †25 ms, ‡peak over 0–50 ms, §late interframe. P: published cost. $*$: SNN event branch only. Ours share a frozen endpoint detector and report complete dense-stream amortized costs.

(b) Event-only detection

Method	AP ↑	AP$_{50}$ ↑
FARSE-CNN [47]	0.4	1.3
DAGR-N [11]	16.2	34.7
DAGR-S [11]	17.8	36.2
DAGR-M [11]	18.8	38.0
DAGR-L [11]	18.9	38.1
SSLA-S [18]	22.9	49.6
SSLA-B [18]	23.4	50.6
SSLA-M [18]	26.2	54.6
SSLA-L [18]	26.8	55.2
Ours S	33.01	58.38
Ours B	33.06	58.41
Ours M	<u>35.56</u>	<u>62.62</u>
Ours L	**36.08**	**63.21**

and DAGR-R50 use common validation at approximately 180 Hz. P marks published costs, and $*$ marks costs of EFH's SNN event branch only. SpikeYOLO's per-event costs are unreported and remain blank.

DSEC event-only detection. We train four event-only models at 430×640 resolution, with temporal SSMs processing the current packet and two preceding 50 ms packets. We choose checkpoints and test-time augmentation (TTA) on validation, then evaluate 12,131 held-out test frames. All four sizes exceed reported SSLA-L in both metrics (Table 2(b)). L reaches 36.08 AP and 63.21 AP50, gains of 9.28 and 8.01 points over SSLA-L, and 17.18 and 25.11 points over event-only DAGR-L [18].

4.3. Efficiency and model scaling

For DSEC and N-Caltech101, complete-input cost includes event encoding, detection, memory updates and final box processing. We divide total work by unique filtered events. Repeated queries add computation without counting their events twice. DSEC covers 526.11 seconds and 833.33 million events at approximately 192 queries per second. N-Caltech101 covers 906 clips and 44.46 million events.

On DSEC, L requires 0.56 MFLOPs/event and 0.45 μJ/event. These are 11.48 and 12.02 times lower than regular DAGR's published costs, and below directed DAGR's 0.57 MFLOPs/event and 0.96 μJ/event. Choosing S saves another 29.12% of computation but loses 11.16 interframe AP points. The shared endpoint detector retains its accuracy and cost.

On N-Caltech101, M reaches 83.79 AP50 at 0.05 MFLOPs/event and 0.04 μJ/event (Fig. 3). It exceeds SSLA-L by 9.49 AP50 points with 17.15 times less computation and 18.52 times less MAC-based energy.

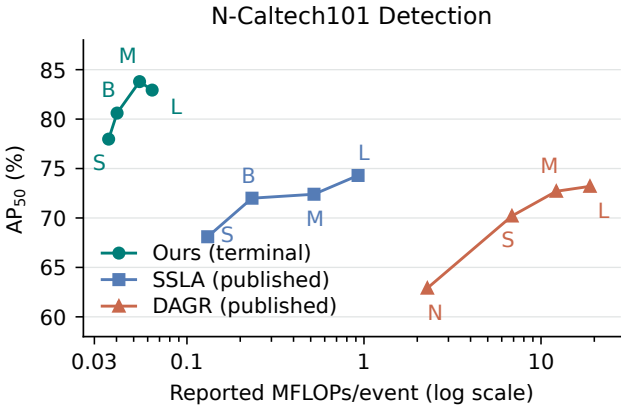


Figure 3. **Accuracy and computation on N-Caltech101 Detection.** Each point represents one model size. Our models combine higher AP50 with lower computation per event. Our costs cover one complete prediction per clip, while SSLA and DAGR retain their published incremental costs.

Figure 4 compares peak GPU memory on Gen1 and N-Caltech101. Across corresponding sizes, SSLA uses 2.06–4.09 times our inference memory and 2.92–4.69 times our training memory. Our inference memory is also lower than all four DAGR configurations. Training memory is closer, with our L slightly higher than DAGR-L on both datasets.

These comparisons use the accounting scopes stated in the tables. Gen1 reports the neural cost of one packet divided by its historical mean of 49,291 events, excluding representation work and unsupported operations. Energy is computed from MACs as defined in Sec. 3.4, rather than measured at the

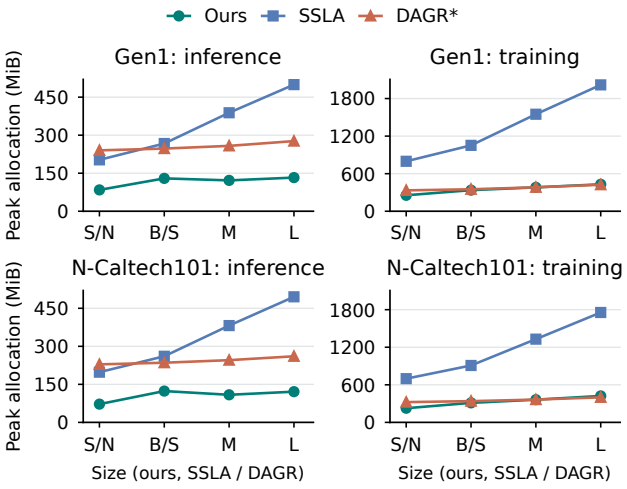


Figure 4. **GPU memory on Gen1 and N-Caltech101 Detection.** Peak FP32 tensor allocation at batch size one on eight shared clips per dataset, capped at 50k events. Training includes gradients and AdamW, inference includes decoding and NMS. Sizes are S/B/M/L for ours and SSLA, N/S/M/L for DAGR. ∗DAGR uses dataset-adapted released architectures with random weights. These Python packet profiles exclude CPU preprocessing, teachers and EMA, and do not measure incremental deployment.

Table 3. **What matters on Gen1?** Validation AP for model M (mean ± SD over three seeds) and paired changes from the control. These seven configurations were confirmed after screening.

Change from control	AP ± SD	ΔAP
Full control	40.17 ± 0.14	0.00
Direct RGB initialization	40.27 ± 0.10	$+0.10$
Training from scratch	33.15 ± 0.19	-7.02
RetinaNet [33] package	39.01 ± 0.27	-1.16
No crop/translation	37.70 ± 0.44	-2.47
Mask average polarity	38.59 ± 0.36	-1.58
Add recent-event context	$\mathbf{40.71 \pm 0.39}$	$+0.54$

GPU.

4.4. Ablation studies

Table 3 tests which training and detection choices help our M-sized SSM on Gen1 validation. The full recipe reaches 40.17 AP. Each variant uses the same training limit and stopping rule.

The control starts from weights pretrained on N-ImageNet classification. Starting from random weights instead loses 7.02 AP, showing the value of pretraining at this budget. Starting directly from RGB-pretrained weights changes AP by only $+0.10$. The uncertainty interval includes zero, so these runs do not establish an advantage for event pretraining over RGB.

Removing random crops and translations loses 2.47 AP. To test polarity, we zero its channels during training and evaluation while preserving event times, counts and input dimensions. The 1.58 AP loss shows that the direction of brightness changes supplies useful information. Replacing the detection path with RetinaNet loses 1.16 AP. Its pyramid, head, box matching and losses change together, while retaining the same SSM backbone architecture.

Finally, we append four statistics from up to 10,000 events in the last 100 ms. A 1×1 projection maps the resulting 24 channels back to 20 before the same backbone. This emphasizes recent activity already in the packet and adds 0.54 AP at 2.08% more counted network computation, excluding event encoding. Resampling validation recordings gives a 95% interval of $[0.04, 0.95]$ AP for the gain in these trained models. This validation variant does not replace our reported test model.

Table 4. **Event classification.** Top-1/Top-5 accuracy (%) against selected published references with their original protocols. Ours uses a 1,742-example N-Caltech101 holdout and three 256px temporal views on official N-ImageNet validation. Bold marks the best listed score per dataset. Dashes indicate unreported values.

Method	Top-1	Top-5
N-Caltech101		
EST [12]	81.70	–
E2VID [44]	86.60	–
ECDP [59]	87.66	–
MEM-NImNet [24]	90.10	–
Video-to-Events [13]	90.60	–
Ours	**91.91**	**97.70**
N-ImageNet		
MEM [24]	57.89	–
PASS-M-K(2) [65]	61.32	–
EventBind (ViT-L/14) [64]	63.54	–
ECDP [59]	64.83	86.30
STP [32]	68.87	**89.65**
Ours	**69.68**	88.98

4.5. Event classification

Table 4 tests whether the architecture also supports recognition without box prediction. N-Caltech101 reaches 91.91% top-1 on the complete 1,742-example holdout, exceeding the reported MEM-NImNet result by 1.81 points. The five-bin model outperforms its ten-bin counterpart by 0.52 points, so finer temporal partitioning is not uniformly helpful. This classification split contains 101 classes and is separate from the detection split.

On the official 50,000-example N-ImageNet validation set, three temporal views through one shared model yield 69.68% top-1 and 88.98% top-5. Top-1 exceeds PASS-M-K(2), EventBind-L and STP by 8.36, 6.14 and 0.81 points, respectively. STP retains higher top-5.

5. Conclusion

We show that event packets can support accurate object detection at low computational cost through compact SSM hidden states. The mixer's main channel transformations act

on a small set of states, while reconstructed spatial features provide the detail needed to locate objects. For RGB–event detection, a cheaper branch reuses image features and saved temporal state between full predictions at image arrivals. Across Gen1, N-Caltech101 Detection and DSEC, our detectors improve accuracy over strong asynchronous baselines. On DSEC and N-Caltech101 Detection, these gains accompany lower amortized computation and MAC energy. These results demonstrate that packet-based spatial context can be retained within an efficient SSM detector.

LLM Usage Statement

We used GPT-6 Astra to assist with code implementation and manuscript revision, and to obtain simulated reviewer feedback. The authors checked and revised all code contributed by the model for correctness and consistency with the proposed method. The research ideas, experimental design, and analysis originated with the authors. References were added manually without LLM assistance. The authors take full responsibility for the work.

Acknowledgments

This work was supported by the European Union's Horizon Europe Research and Innovation Programme under grant agreement No. 101120732 (AUTOASSESS), the European Research Council (ERC) under grant agreement No. 864042 (AGILEFLIGHT), and the National University of Singapore.

Compact SSM Hidden States for Accurate and Efficient Event-Based Vision

Supplementary Material

6. Event-only semantic segmentation on DSEC

Detection needs only a few boxes per image. Semantic segmentation [10] is harder: it must assign a class to every pixel. It therefore tests whether the compact SSM hidden states of Sec. 3.1 keep enough spatial detail for dense prediction. Each mixer reconstructs features at all N locations, so a standard dense decoder can be placed on top of the same backbone.

Task and data. We use DSEC-Semantic [51]. Each pixel belongs to one of 11 classes: background, building, fence, person, pole, road, sidewalk, vegetation, car, wall, and traffic sign. Labels have a resolution of 640×440, with 8,082 training and 2,809 test frames. The model sees only events at inference.

Model. For each prediction, we accumulate events over 25, 100 and 300 ms. Short windows show fast motion, while long windows also show static structure. Five temporal bins per window give 15 input channels. Our compact SSM backbone turns this tensor into multiscale spatial features. A causal temporal SSM combines the features of four successive event tensors, so each prediction uses only current and past events. A UPerNet decoder [56] then predicts a class score for every pixel.

Training. First, frozen DINOv3 image features [49] teach an event encoder to produce similar features from paired events. This uses 109,645 image–event samples from Cityscapes [6], KITTI [16], DECD [45], and VisEvent/COESOT [52, 55]. This encoder initializes a semantic event teacher, which we pretrain together with our compact SSM student for 31,250 updates. On DSEC, the student learns from the labels and from the teacher's soft class predictions. The distillation weight is 1 for the first 12,000 updates and 4 until 24,000. The final model averages the last three checkpoints and keeps the batch-normalization statistics of the last one. Training uses 16 GH200 GPUs, BF16 precision, and a global batch of 32.

Selection without test data. We fit one additive score offset per class on the Zurich training sequences `00_a` and `08_a` (1,726 frames). We confirm the recipe on a second holdout, `01_a` and `07_a` (2,144 frames). Each fold's teacher is trained without that fold. We then retrain on all 8,082 training frames and evaluate each seed once on the official test split. Seed 0 was chosen as the headline result before evaluation. Each seed uses one native 640×440 view, with no flipping or other test-time augmentation.

Table 5. **Our three seeds on the DSEC-Semantic test split.** Mean $\pm$ population SD.

Run	mIoU (%) ↑	Pixel acc. (%) ↑
Seed 0 (prespecified)	70.08	93.77
Seed 1	69.89	93.78
Seed 2	70.10	93.77
Mean $\pm$ SD	$\mathbf{70.02 \pm 0.10}$	93.77 ± 0.01

Metrics. For each class c, the intersection over union compares the pixels predicted as c with the pixels labelled as c:

$$\mathrm{IoU}_c = \frac{\mathrm{TP}_c}{\mathrm{TP}_c + \mathrm{FP}_c + \mathrm{FN}_c},$$

where TP_c, FP_c and FN_c count correct, wrongly added, and missed pixels. The mean IoU (mIoU) averages IoU_c over all 11 classes, so every class counts equally. Pixel accuracy is the fraction of all labelled pixels whose class is predicted correctly. Large classes such as road, building and background dominate it, so a model can score high while missing thin classes such as poles and signs. We accumulate counts over all 2,809 test frames before computing both metrics. We compare with other methods only in mIoU, because it is the one metric that every paper reports with the same definition. Accuracy scores in prior work do not share one definition: ScaleEvent reports average class accuracy [5], and other papers do not always state which accuracy they use. We therefore report our pixel accuracy only for completeness.

Results. Tab. 5 shows that the three seeds reach 70.02 ± 0.10 mIoU, so the result is stable across seeds. Tab. 6 compares this with published results. Our mean is 0.37 points above ScaleEvent-L, the strongest event-only method we found, and 0.98 points above Hyper-RED-L. Both use a ViT-L/16 encoder distilled from DINOv3. The margin is small, and training data and supervision differ across papers. We therefore read this result as matching the strongest event-only methods, not as a clear improvement over them. The class-level results show where the model differs. Fence IoU averages 36.81%, against 26.88% for ScaleEvent-L. Wall IoU averages 45.25%, against 52.53%.

Table 6. **Comparison on DSEC-Semantic.** Event-only methods, mIoU (%) as reported by each paper. Training data, supervision, and test-time augmentation differ across papers.

Method	Backbone	mIoU ↑
ESS (supervised) [51]	E2VID	53.29
OpenESS [25]	E2VID	57.21
ECDP [59]	ResNet-50	59.16
ECDDP [60]	Swin-T	61.25
STP [32]	Swin-T	62.05
TESPEC [38]	recurrent Swin-T	62.77
GEP [2]	ViT-B	67.37
Hyper-RED [54]	ViT-L/16	69.04
ScaleEvent [5]	ViT-L/16	69.65
Ours (3-seed mean)	compact SSM	**70.02**

References

[1] Marco Cannici, Marco Ciccone, Andrea Romanoni, and Matteo Matteucci. A differentiable recurrent surface for asynchronous event-based data. In *European Conference on Computer Vision*, pages 136–152. Springer, 2020. 6

[2] Jianwen Cao, Jiaxu Xing, Nico Messikommer, and Davide Scaramuzza. Generative event pretraining with foundation model alignment. *arXiv preprint arXiv:2603.23032*, 2026. 2

[3] Haosheng Chen, Lian Luo, Mengjingcheng Mo, Zhanjie Wu, Ji Gan, Jiaxu Leng, and Xinbo Gao. Ehgcn: Hierarchical euclidean-hyperbolic fusion via motion-aware gcn for hybrid event stream perception. *arXiv preprint arXiv:2504.16616*, 2025. 6

[4] Nicholas FY Chen. Pseudo-labels for supervised learning on dynamic vision sensor data, applied to object detection under ego-motion. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 757–75709. IEEE, 2018. 6

[5] Zhiwen Chen, Junhui Hou, Zhiyu Zhu, Jinjian Wu, and Guangming Shi. Scaling dense event-stream pretraining from visual foundation models. *arXiv preprint arXiv:2603.03969*, 2026. 1, 2

[6] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 1

[7] Tri Dao and Albert Gu. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. *Proceedings of the 41st International Conference on Machine Learning*, 2024. 2, 3, 4

[8] Pierre De Tournemire, Davide Nitti, Etienne Perot, Davide Migliore, and Amos Sironi. A large scale event-based detection dataset for automotive. *arXiv preprint arXiv:2001.08499*, 2020. 5

[9] Rui Fan, Weidong Hao, Juntao Guan, Lai Rui, Lin Gu, Tong Wu, Fanhong Zeng, and Zhangming Zhu. Eventpillars: Pillar-based efficient representations for event data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2861–2869, 2025. 2, 3, 6

[10] Sergei Fedorchenko, Max Zehnder, Nikola Zubić, Rong Zou, and Davide Scaramuzza. Event-SAMA: Language-grounded video object segmentation on neuromorphic event streams. In *Conference on Neural Information Processing Systems Workshops (NeurIPSW)*, 2026. 1

[11] Daniel Gehrig and Davide Scaramuzza. Low-latency automotive vision with event cameras. *Nature*, 629(8014):1034–1040, 2024. 1, 2, 3, 6, 7

[12] Daniel Gehrig, Antonio Loquercio, Konstantinos G Derpanis, and Davide Scaramuzza. End-to-end learning of representations for asynchronous event-based data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5633–5643, 2019. 2, 8

[13] Daniel Gehrig, Mathias Gehrig, Javier Hidalgo-Carrió, and Davide Scaramuzza. Video to events: Recycling video datasets for event cameras. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3583–3592. IEEE, 2020. 8

[14] Mathias Gehrig and Davide Scaramuzza. Recurrent vision transformers for object detection with event cameras. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13884–13893. IEEE, 2023. 2

[15] Mathias Gehrig, Willem Aarents, Daniel Gehrig, and Davide Scaramuzza. Dsec: A stereo event camera dataset for driving scenarios. *IEEE Robotics and Automation Letters*, 6(3):4947–4954, 2021. 2

[16] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 32(11):1231–1237, 2013. 1

[17] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In *First Conference on Language Modeling*, 2024. 3

[18] Haiqing Hao, Zhipeng Sui, Rong Zou, Zijia Dai, Nikola Zubić, Davide Scaramuzza, and Wenhui Wang. Low-latency event-based object detection with spatially-sparse linear attention. In *European Conference on Computer Vision*, pages 20–38. Springer, 2026. 2, 3, 6, 7

[19] Haiqing Hao, Nikola Zubic, Weihua He, Zhipeng Sui, Davide Scaramuzza, and Wenhui Wang. Maximizing asynchronicity in event-based neural networks. In *International Conference on Learning Representations*, pages 123834–123855, 2026. 6

[20] Massimiliano Iacono, Stefan Weber, Arren Glover, and Chiara Bartolozzi. Towards event-driven object detection with off-the-shelf deep learning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–9. IEEE, 2018. 6

[21] Zhuangyi Jiang, Pengfei Xia, Kai Huang, Walter Stechele, Guang Chen, Zhenshan Bing, and Alois Knoll. Mixed frame-/event-driven fast pedestrian detection. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8332–8338. IEEE, 2019. 6

[22] Kirill S Khitushkin, Tim T Isakov, and Aleksandr V Bakhshiev. Using spiking neural networks for event and

multimodal data processing in object detection tasks. In *International Conference on Neuroinformatics*, pages 123–137. Springer, 2025. 6, 7

[23] Junho Kim, Jaehyeok Bae, Gangin Park, Dongsu Zhang, and Young Min Kim. N-imagenet: Towards robust, fine-grained object recognition with event cameras. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2126–2136. IEEE, 2021. 3

[24] Simon Klenk, David Bonello, Lukas Koestler, Nikita Araslanov, and Daniel Cremers. Masked event modeling: Self-supervised pretraining for event cameras. In *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2367–2377. IEEE, 2024. 8

[25] Lingdong Kong, Youquan Liu, Lai Xing Ng, Benoit R Cottereau, and Wei Tsang Ooi. Openess: Event-based semantic scene understanding with open vocabularies. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15686–15698. IEEE, 2024. 2

[26] Sanghyeok Lee, Joonmyung Choi, and Hyunwoo J Kim. Efficientvim: Efficient vision mamba with hidden state mixer based state space duality. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14923–14933. IEEE, 2025. 2, 3, 4

[27] Dianze Li, Jianing Li, Xu Liu, Xiaopeng Fan, and Yonghong Tian. Asynchronous collaborative graph representation for frames and events. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1655–1666. IEEE, 2025. 3

[28] Dianze Li, Jianing Li, Xu Liu, Zhaokun Zhou, Xiaopeng Fan, and Yonghong Tian. Rethinking neuromorphic object detection with hybrid dynamic interaction transformers: D. li et al. *International Journal of Computer Vision*, 134(7):340, 2026. 3

[29] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *Advances in neural information processing systems*, 33:21002–21012, 2020. 5

[30] Yijin Li, Han Zhou, Bangbang Yang, Ye Zhang, Zhaopeng Cui, Hujun Bao, and Guofeng Zhang. Graph-based asynchronous event processing for rapid object recognition. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 914–923. IEEE, 2021. 6

[31] Ziqi Li, Tao Gao, Yisheng An, Ting Chen, Jing Zhang, Yuanbo Wen, Mengkun Liu, and Qianxi Zhang. Brain-inspired spiking neural networks for energy-efficient object detection. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3552–3562. IEEE, 2025. 3

[32] Quanmin Liang, Qiang Li, Shuai Liu, Xinzi Cao, Jinyi Lu, Feidiao Yang, Wei Zhang, Kai Huang, and Yonghong Tian. Efficient event camera data pretraining with adaptive prompt fusion. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8656–8667. IEEE, 2025. 3, 8, 2

[33] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 8

[34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2017. 5

[35] Dongyue Lu, Lingdong Kong, Gim Hee Lee, Camille Simon Chane, and Wei Ooi. Flexevent: towards flexible event-frame object detection at varying operational frequencies. *Advances in Neural Information Processing Systems*, 38:89009–89044, 2026. 3

[36] Yunfan Lu, Nico Messikommer, Xiaogang Xu, Liming Chen, Yuhan Chen, Nikola Zubić, Davide Scaramuzza, and Hui Xiong. Hybrid event–frame sensors: Modeling, calibration, and simulation. In *European Conference on Computer Vision*, pages 167–187. Springer, 2026. 3

[37] Nico Messikommer, Daniel Gehrig, Antonio Loquercio, and Davide Scaramuzza. Event-based asynchronous sparse convolutional networks. In *European Conference on Computer Vision*, pages 415–431. Springer, 2020. 6

[38] Mohammad Mohammadi, Ziyi Wu, and Igor Gilitschenski. Tespec: Temporally-enhanced self-supervised pretraining for event cameras. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7782–7793. IEEE, 2025. 2

[39] Garrick Orchard, Ajinkya Jayawant, Gregory K Cohen, and Nitish Thakor. Converting static image datasets to spiking neuromorphic datasets using saccades. *Frontiers in neuroscience*, 9:437, 2015. 5

[40] Yansong Peng, Yueyi Zhang, Zhiwei Xiong, Xiaoyan Sun, and Feng Wu. Get: Group event transformer for event-based vision. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6015–6025. IEEE, 2023. 6

[41] Yansong Peng, Hebei Li, Yueyi Zhang, Xiaoyan Sun, and Feng Wu. Scene adaptive sparse transformer for event-based object detection. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16794–16804. IEEE, 2024. 6

[42] Etienne Perot, Pierre De Tournemire, Davide Nitti, Jonathan Masci, and Amos Sironi. Learning to detect objects with a 1 megapixel event camera. *Advances in Neural Information Processing Systems*, 33:16639–16652, 2020. 2, 6

[43] Gokul Raju Govinda Raju, Nikola Zubic, Marco Cannici, and Davide Scaramuzza. Perturbed state space feature encoders for optical flow with event cameras. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 5072–5081. IEEE, 2025. 3

[44] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. Events-to-video: Bringing modern computer vision to event cameras. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3852–3861. IEEE, 2019. 8

[45] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. High speed and high dynamic range video with an event camera. *IEEE transactions on pattern analysis and machine intelligence*, 43(6):1964–1980, 2019. 1

[46] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian D. Reid, and Silvio Savarese. Generalized

intersection over union: A metric and a loss for bounding box regression. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 658–666, 2019. 5

[47] Riccardo Santambrogio, Marco Cannici, and Matteo Matteucci. Farse-cnn: Fully asynchronous, recurrent and sparse event-based cnn. In *European conference on computer vision*, pages 1–18. Springer, 2024. 3, 6, 7

[48] Simon Schaefer, Daniel Gehrig, and Davide Scaramuzza. Aegnn: Asynchronous event-based graph neural networks. In *2022 IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 12361–12371. IEEE, 2022. 2, 3, 6

[49] Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025. 1

[50] Taylan Soydan, Nikola Zubić, Nico Messikommer, Siddhartha Mishra, and Davide Scaramuzza. S7: Selective and simplified state space layers for sequence modeling. *arXiv preprint arXiv:2410.03464*, 2024. 3

[51] Zhaoning Sun, Nico Messikommer, Daniel Gehrig, and Davide Scaramuzza. Ess: Learning event-based semantic segmentation from still images. In *European Conference on Computer Vision*, pages 341–357. Springer, 2022. 1, 2

[52] Chuanming Tang, Xiao Wang, Ju Huang, Bo Jiang, Lin Zhu, Shifeng Chen, Jianlin Zhang, Yaowei Wang, and Yonghong Tian. Revisiting color-event based tracking: A unified network, dataset, and metric. *Pattern Recognition*, page 112718, 2025. 1

[53] Jiabao Wang, Yuming Chen, Zhaohui Zheng, Xiang Li, Ming-Ming Cheng, and Qibin Hou. Crosskd: Cross-head knowledge distillation for object detection. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16520–16530. IEEE, 2024. 5

[54] Meisen Wang, Zhiqiang Tian, Wei Bao, Chengjie Wang, Shaoyi Du, and Siqi Li. Hyper-red: Scalable event pretraining via semantic hypergraph distillation. *arXiv preprint arXiv:2609.16811*, 2026. 2

[55] Xiao Wang, Jianing Li, Lin Zhu, Zhipeng Zhang, Zhe Chen, Xin Li, Yaowei Wang, Yonghong Tian, and Feng Wu. VisEvent: Reliable object tracking via collaboration of frame and event flows. *IEEE Transactions on Cybernetics*, 2023. 1

[56] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *European conference on computer vision*, pages 432–448. Springer, 2018. 1

[57] Nan Yang, Yang Wang, Zhanwen Liu, Meng Li, Yisheng An, and Xiangmo Zhao. Smamba: Sparse mamba for event-based object detection. In *Proceedings of the AAAI conference on artificial intelligence*, pages 9229–9237, 2025. 3, 6

[58] Nan Yang, Yang Wang, Zhanwen Liu, Yuchao Dai, Yang Liu, and Xiangmo Zhao. Focus through motion: Rgb-event collaborative token sparsification for efficient object detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 2026. 3

[59] Yan Yang, Liyuan Pan, and Liu Liu. Event camera data pretraining. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10665–10675. IEEE, 2023. 8, 2

[60] Yan Yang, Liyuan Pan, and Liu Liu. Event camera data dense pre-training. In *European Conference on Computer Vision*, pages 292–310. Springer, 2024. 2

[61] Chengjun Zhang, Yuhao Zhang, Jisong Yu, Jie Yang, and Mohamad Sawan. Event-fused hybrid ann-snn architecture for low-latency object detection in automotive vision. *IEEE Robotics and Automation Letters*, 11(3):3622–3628, 2026. 6, 7

[62] Haoyang Zhang, Ying Wang, Feras Dayoub, and Niko Sunderhauf. Varifocalnet: An iou-aware dense object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8514–8523, 2021. 5

[63] Haitian Zhang, Xiangyuan Wang, Chang Xu, Xinya Wang, Fang Xu, Huai Yu, Lei Yu, and Wen Yang. Frequency-adaptive low-latency object detection using events and frames. *arXiv preprint arXiv:2412.04149*, 2024. 3

[64] Jiazhou Zhou, Xu Zheng, Yuanhuiyi Lyu, and Lin Wang. Eventbind: Learning a unified representation to bind them all for event-based open-world understanding. In *European Conference on Computer Vision*, pages 477–494. Springer, 2024. 8

[65] Jiazhou Zhou, Kanghao Chen, Lei Zhang, and Lin Wang. Pass: path-selective state space model for event-based recognition. *Advances in Neural Information Processing Systems*, 38:168236–168264, 2026. 8

[66] Liang-Hui Zhu, Ben-Cheng Liao, Qian Zhang, Xin-Long Wang, Wen-Yu Liu, and Xing-Gang Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model. In *International Conference on Machine Learning*, 2024. 3

[67] Nikola Zubic and Davide Scaramuzza. Gg-ssms: Graph-generating state space models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 28863–28873, 2025. 3

[68] Nikola Zubić and Davide Scaramuzza. Regularity and stability properties of selective ssms with discontinuous gating. *Transactions on Machine Learning Research*, 2026. 3

[69] Nikola Zubić, Daniel Gehrig, Mathias Gehrig, and Davide Scaramuzza. From chaos comes order: Ordering event representations for object recognition and detection. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12800–12810. IEEE, 2023. 2, 6

[70] Nikola Zubic, Mathias Gehrig, and Davide Scaramuzza. State space models for event cameras. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5819–5828, 2024. 3, 6

[71] Nikola Zubić, Qian Li, Yuyi Wang, and Davide Scaramuzza. On the expressive power and limitations of multi-layer ssms. *Advances in Neural Information Processing Systems*, 2026. 3