

Spiking Neural Network as Adaptive Event Stream Slicer

Jiahang Cao^{1*} Mingyuan Sun^{2*} Ziqing Wang^{3*} Hao Cheng¹
Qiang Zhang^{1,4} Shibo Zhou⁵ Renjing Xu¹

¹ The Hong Kong University of Science and Technology (Guangzhou)

² Northeastern University ³ Northwestern University

⁴ Beijing Innovation Center of Humanoid Robotics Co. Ltd. ⁵ Brain Mind Innovation
{jcao248, hcheng046, qzhang749}@connect.hkust-gz.edu.cn
mingyuansun20@gmail.com, ziqingwang2029@u.northwestern.edu
bob@brain-mind.com.cn, renjingxu@hkust-gz.edu.cn

Abstract

Event-based cameras are attracting significant interest as they provide rich edge information, high dynamic range, and high temporal resolution. Many state-of-the-art event-based algorithms rely on splitting the events into fixed groups, resulting in the omission of crucial temporal information, particularly when dealing with diverse motion scenarios (*e.g.*, high/low speed). In this work, we propose **SpikeSlicer**, a novel-designed plug-and-play event processing method capable of splitting events stream adaptively. SpikeSlicer utilizes a low-energy spiking neural network (SNN) to trigger event slicing. To guide the SNN to fire spikes at optimal time steps, we propose the Spiking Position-aware Loss (SPA-Loss) to modulate the neuron’s state. Additionally, we develop a Feedback-Update training strategy that refines the slicing decisions using feedback from the downstream artificial neural network (ANN). Extensive experiments demonstrate that our method yields significant performance improvements in event-based object tracking and recognition. Notably, SpikeSlicer provides a brand-new SNN-ANN cooperation paradigm, where the SNN acts as an efficient, low-energy data processor to assist the ANN in improving downstream performance, injecting new perspectives and potential avenues of exploration.

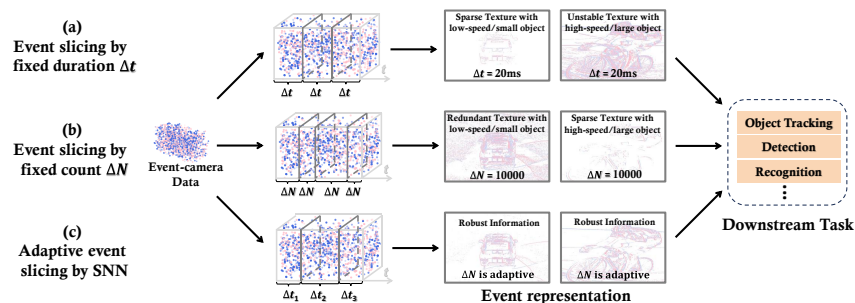


Figure 1: Comparison of event slicing methods. Traditional methods slice event streams based on prefixed time intervals (a) or event counts (b). In contrast, our approach (c) utilizes SNN as a dynamic event processor for adaptive event slicing. The sliced sub-event streams can be converted into various event representations with robust information and then applied to multiple downstream tasks.

*Equal contribution.

1 Introduction

Event-based cameras [1] are bio-inspired sensors that capture event streams in an asynchronous and sparse way. Compared with conventional frame-based cameras, event-based cameras offer numerous outstanding properties: high temporal resolution (with the order of μs), high dynamic range (higher than 120 dB), low latency, and low power consumption. Over recent years, rapid growth has been witnessed in dealing with event data due to the inherent advantages of event-based cameras, such as object tracking [2, 3], depth estimation [4, 5], and recognition [6, 7]. Before applying to various downstream tasks, the event stream must be split by groups and then transformed into different event representations, *e.g.*, frame, voxel, or point for deep learning architectures.

In detail, the process of event-to-representation conversion consists of two steps: (1) slicing the raw event stream into multiple sub-event stream groups, and (2) converting these sub-event streams into different event representations. Much of the current research focuses on the second step, aiming to refine event representation [8] techniques such as time surface [6] and event spike tensor [9]. Yet, this focus often overlooks the crucial first step of slicing, where issues such as the non-uniformity of the information contained in the fixed-sliced event remain unaddressed.

To address the challenges in the slicing process, we delve into the limitations of traditional slicing techniques. Common methods typically cut the event stream into several fixed groups. For example, slicing event stream with fixed event count [10] or fixed time intervals [11, 12] as depicted in Fig. 1. However, these fixed-group slicing techniques often lead to problems: they may cause insufficient information capture in low-speed motion scenarios or excessive redundancy in high-speed conditions, thereby failing to accurately capture the dynamic changes in event distribution. Additionally, some hyper-parameters, *e.g.*, the length of time interval, are highly-sensitive to the downstream tasks (examples are provided in Appendix C) and must be carefully pre-determined. Although some latest slicing methods [13, 14] propose to adaptively sample the events, there still exists the problem of hyper-parameter tuning which can not achieve a fully learnable and adaptable slicing process.

In order to address the above issues, we propose SpikeSlicer, a novel-design event processing method that can adaptively slice the event streams. To achieve this, SpikeSlicer utilizes an SNN as an event trigger to dynamically determine the optimal moment to split the event stream. Our objectives include: (1) training the SNN to spike at a specific time step for accurate event slicing, and (2) developing a training strategy to identify the best slicing time for a continuous event stream during training. In our paper, we achieve (1) through our newly introduced Spiking Position-aware Loss (SPA-Loss) function, which effectively guides the SNN to spike at the desired time by manipulating the membrane potential. For (2), we implement a Feedback-Update training strategy, where the SNN receives real-time performance feedback from the downstream ANN model for supervision. An overview of our proposed method is depicted in Fig. 2. We evaluate the effectiveness of our proposed SpikeSlicer in two downstream tasks: (i) event-based object tracking, which is strongly sensitive to temporal information and motion dynamics, and (ii) event-based recognition, which is highly related to event density. Extensive experiments validate the effectiveness of the proposed approach.

To sum up, our contributions are as follows:

- We propose SpikeSlicer, a novel plug-and-play event processing method capable of splitting event streams in an adaptive manner.
- We design the SPA-Loss to guide the SNN to trigger spikes at the expected time steps. We then propose a novel Feedback-Update strategy that optimizes the event slicing process based on the ANN feedback.
- Extensive experiments demonstrate that SpikeSlicer significantly improves model performance by up to 22.3% in object tracking and 19.2% in recognition with only a 0.32% increase in energy consumption.

2 Background and Related Work

Event-based Cameras. They are bio-inspired sensors, which capture the relative intensity changes asynchronously. In contrast to standard cameras that output 2D images, event cameras output sparse event streams. When brightness change exceeds a threshold C , an event e_k is generated containing position $\mathbf{u} = (x, y)$, time t_k , and polarity p_k : $\Delta L(\mathbf{u}, t_k) = L(\mathbf{u}, t_k) - L(\mathbf{u}, t_k - \Delta t_k) = p_k C$. The polarity of an event reflects the direction of the changes (*i.e.*, brightness increase (“ON”) or

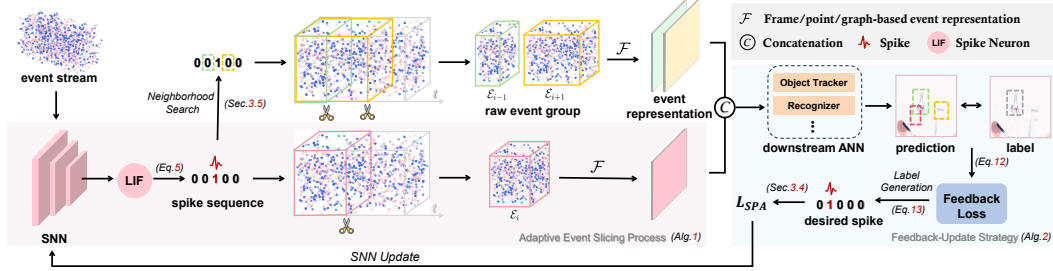


Figure 2: Overview of our method. The input events are first fed into an SNN, and the event is determined to be sliced when a spike occurs. To find the accurate slicing time, the neighborhood search method explores other time steps, and feeds event representations to the downstream ANN model (e.g., object tracker or recognizer). The ANN model then offers feedback, which guides the SNN in firing spikes at the optimal slicing time by supervising the membrane potential through the Spiking Position-aware Loss \mathcal{L}_{SPA} .

decrease (“OFF”). In general, the output of an event camera is a sequence of events, which can be described as: $\mathcal{E} = \{e_k\}_{k=1}^N = \{\{\mathbf{u}_k, t_k, p_k\}\}_{k=1}^N$. With the advantages of high temporal resolution, high dynamic range, and low energy consumption, event cameras are gradually attracting attention in the fields of tracking [3, 15], identification [7], 3D reconstruction [16, 17] and estimation [5].

Spiking Neural Network (SNN). SNNs are potential competitors to artificial neural networks (ANNs) due to their distinguished properties: high biological plausibility, event-driven nature, and low power consumption. In SNNs, all information is represented by binary time series data rather than float representation, leading to energy efficiency gains. Also, SNNs possess powerful abilities to extract spatial-temporal features for various tasks, including recognition [18], tracking [2], and image generation [19]. In this paper, we adopt the widely used Leaky Integrate-and-Fire (LIF, [20, 21]) model, which effectively characterizes the dynamic process of spike generation and can be defined as:

$$V[n] = \beta V[n-1] + \gamma I[n], \quad (1)$$

$$S[n] = \Theta(V[n] - \vartheta_{th}), \quad (2)$$

where n is the time step and β is the leaky factor that controls the information reserved from the previous time step; $V[n]$ is the membrane potential; $S[n]$ denotes the output spike which equals 1 when there is a spike and 0 otherwise; $\Theta(x)$ is the Heaviside function. When the membrane potential exceeds the threshold ϑ_{th} , the neuron will trigger a spike and resets its membrane potential to $V_{reset} < \vartheta_{th}$. Meanwhile, when $\beta = \gamma = 1$, LIF neuron evolves into Integrate-and-Fire (IF) neuron. We also introduce a no-reset membrane potential $U[n]$, meaning that the membrane potential does not reset, but directly passes the original value to the next time step (i.e., $U[n] = V[n]$ after Eq. 2).

3 Our Approach: SpikeSlicer

In this section, we first introduce the concept of event cells for data preparation (Sec. 3.1). We then introduce the adaptive event slicing process by utilizing an SNN as the event trigger (Sec. 3.2). In Sec. 3.3, we introduce a novel Spiking Position-aware Loss (SPA-Loss) to supervise the SNN to slice the event at the precise time. Lastly, we build a feedback-update (Sec. 3.4) strategy that allows the resulting events to be correlated with the feedback from the downstream model, thereby improving overall performance.

3.1 Converting Event Stream to Event Cell

Event streams are asynchronous data that can be represented as a set: $\mathcal{E} = \{[x_i, y_i, t_i, p_i]\}_{i=1}^N$ with a time span of T (i.e., $t_i \in [t_0, t_0 + T]$). We envision an ideal situation where the SNN is utilized to directly process the raw data stream and slice the event. However, in software simulations, the event stream should be converted into event representations to comply with the input requirements of existing deep learning frameworks. In this paper, following the mainstream research, we adopt the voxel-grid representation [22] as the input of SNN. We first introduce the definition of event cell:

Definition 1 (Event cell). Consider a small time interval δt , event cell is a single-grid event representation in the form of: $C_{\pm}(x, y, t_*) = \mathcal{F}_{\text{voxel}}(G_{\pm}(x, y, t, \{t \in [t_*, t_* + \delta t]\}))$, where $\mathcal{F}_{\text{voxel}}$ denotes the voxel grid [22] method to process the raw event groups G_{\pm} (Appendix F) with $t \in [t_*, t_* + \delta t]$ into a grid representation.

A whole event stream can be then represented by a list of N event cells, *i.e.*, $\{C_{\pm}(x, y, t_0), C_{\pm}(x, y, t_0 + \delta t), \dots, C_{\pm}(x, y, t_0 + (N - 1)\delta t)\}$, where $N = T/\delta t$ and each cell corresponds to a discrete time index $n \in \{0, 1, \dots, N - 1\}$. The mapping of discrete time to real event time interval is defined as:

$$f_{\text{time}}(n) = \{t | t \in [t_0 + n\delta t, t_0 + (n + 1)\delta t]\}. \quad (3)$$

e.g., the time interval of $C_{\pm}(x, y, t_0)$ is $f_{\text{time}}(0)$. In the following sections, we abbreviate the event cell as $C[n]$.

Discussion: Distinct from the typical voxel grid, an event cell only contains a brief time interval, *i.e.*, δt is small. At this point, the entire cell sequence appropriately represents the raw event stream, while simultaneously fulfilling the input requisites for the SNN.

3.2 Adaptive Event Slicing Process

Utilizing SNNs on neuromorphic hardware for processing event streams is low-energy and low-latency [23, 24]. Hence, we adopt the Spiking Neural Network as the event stream slicer, aiming for dynamically slicing the event stream to enhance the downstream performance. Incorporating with the SNN, we now describe the adaptive event slicing process:

Considering an event stream \mathcal{E} , we first convert \mathcal{E} into a list of time-continuous event cells. Event cells are then continuously entered into a SNN (structure details are provided in Appendix L) through a loop operation. Through forward propagation, the features of the last hidden layers (h^{L-1}) are finally mapped to a single spiking neuron to activate spikes:

$$S_{\text{out}} = \text{LIF}(\text{SNN}_{\text{FC}}(h^{L-1})). \quad (4)$$

Once the spiking neuron generates a spike (*i.e.*, $S_{\text{out}} = 1$) at current time n_c , it is considered a slicing action. This allows us to obtain the time interval from the end of the previous spike to the current spike. Suppose the previous spike happened at time n_p , the real event time interval is within $T_{\text{group}} = \bigcup_{n=n_p+1}^{n_c} f_{\text{time}}(n) = \{t \in [t_0 + (n_p + 1)\delta t, t_0 + (n_c + 1)\delta t]\}$. Then, the raw event data within this time interval form an event group, which can be converted to any event representation:

$$\mathcal{D}_{n_c} = \mathcal{F}(G_{\pm}(x, y, t, T_{\text{group}})), \quad (5)$$

where \mathcal{F} denotes an event representation method (*e.g.*, frame [22], point [25], graph [26] and surface [6]-based methods). This representation \mathcal{D}_{n_c} is then prepared for the downstream tasks.

Event Slicing Rule: The slice of the event stream is determined by the state (excited/resting) of the SNN’s spiking neuron. Serving as a dynamic event trigger, SNN promptly decides to split events upon spike generation and extracts the precise time interval of the raw event stream. The details of the adaptive event slicing process is shown in Alg. 1.

3.3 Spiking Position-aware Loss

In this section, we propose the Spiking Position-aware Loss (SPA-Loss), which contains two parts: (1) membrane potential-driven loss (Mem-Loss) is used to directly guide the spiking state of the spiking neuron at a specified timestamp, and (2) linear-assuming loss (LA-Loss), which is designed to resolve the dependence phenomenon between neighboring membrane potentials, achieving a more precise spiking time. Moreover, we introduce a (3) dynamic hyperparameter tuning method to avoid the experimental bias caused by the manual setting of hyperparameters.

3.3.1 Membrane Potential-driven Loss

As mentioned in the previous section, the slice position of event is determined upon the spike occurrence. The challenge now lies in directing the SNN to trigger a spike precisely at the optimal position, once the label for this optimal slicing position is provided (in Sec. 3.4).

Consider consecutive event cells as inputs starting from the previous spiking time, suppose we expect SNN to slice the event at n^* , *i.e.*, a spike S_{out} is triggered at n^* . This corresponds to the membrane potential of the spiking neuron needing to reach the threshold V_{th} at n^* , which inspired us to guide the spike time by directly giving the desired membrane potentials. However, membrane potential returns to the resting state immediately after the occurrence of a spike, which may result in inaccurate guidance at later moments (Appendix G). Thus, we choose to supervise the no-reset membrane

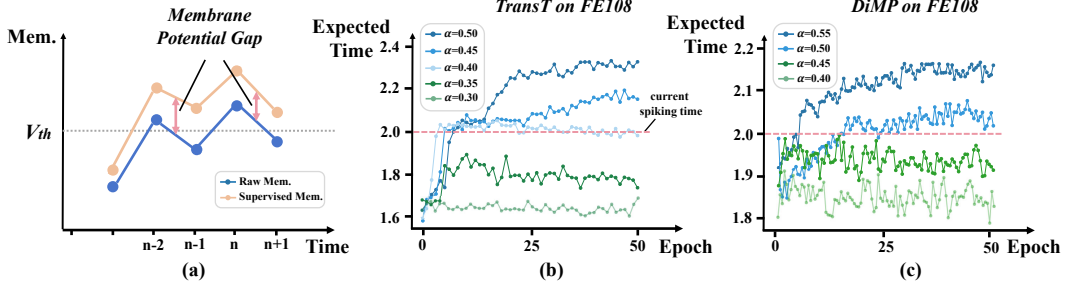


Figure 3: Empirical observations: (a) Hill effect in adaptive slicing process; (b) Impact of hyperparameter α settings on TransT tracker [27] and (c) DiMP tracker [28].

potential $U[n]$ (Eq. 16) to exceed the threshold (i.e., $U[n^*] \geq V_{th}$). The membrane potential-driven loss is defined as:

$$\mathcal{L}_{Mem} = \|U[n^*] - (1 + \alpha)V_{th}\|_2^2, \quad (6)$$

where $\alpha \geq 0$ is a hyperparameter to control the desired membrane potential to exceed the threshold. However, an excessively high α may directly induce a premature spike in the neuron, thereby influencing the membrane potential state at the targeted time step. We provide a proposition to address this problem:

Proposition 1. Suppose the input event cell sequence has length N , desired spiking time is n^* ($n^* \in \{0, 1, \dots, N - 1\}$), the membrane potential at time n^* satisfying the constraints:

$$V_{th} \leq U[n^*] \leq \max(\beta V_{th} + \gamma I[n^*], V_{th}), \quad (7)$$

where $I[n^*]$ is the input synaptic current from Eq. 1. Then the spiking neuron fires a spike at time n^* and does not excite spikes at neighboring moments.

The proof is provided in the Appendix H. Based on the proposition, we modify the loss function into:

$$\mathcal{L}_{Mem} = \|U[n^*] - ((1 - \alpha)U_{lower} + \alpha U_{upper})\|_2^2, \quad (8)$$

where $U_{lower} = V_{th}$ and $U_{upper} = \max(\beta V_{th} + \gamma I[n^*], V_{th})$ denote the lower and upper bounds of the $U[n^*]$ provided in the proposition, respectively; $\alpha \in [0, 1]$ balances the desired membrane potential $U[n^*]$ between U_{lower} and U_{upper} . Experiments in Sec. 4.1 demonstrate that Mem-Loss is able to supervise the SNN to determine the slicing of the event flow at a specified timestamp. More details of Mem-Loss are provided in Appendix I.1.

3.3.2 Linear-assuming Loss

However, only using Mem-Loss is unable to guarantee that the spiking neuron can trigger spikes at any expected timestamp. We have the following observations:

Observation 1 (Hill effect). Suppose there exists a situation where $S[n] = 1$ and $U[n] \geq U[n + 1]$. If the neuron is expected to activate a spike at time $n + 1$, $U[n + 1]$ will be driven to reach the threshold through the Mem-Loss. Nonetheless, the supervised neuron still exhibits $U[n] \geq U[n + 1]$, causing an early spike at time n .

As illustrated in Fig. 3(a), if $U[n] \geq U[n + 1]$ exists, this membrane potential gap is still inherited after the supervision. In this case, when the later membrane potential is guided to exceed the threshold, the earlier membrane potential reaches the threshold sooner and turns the neuron into the resting state. This poses challenges in obtaining a second spike at the later moment.

Therefore, we expect the later membrane potential to increase monotonically with the time step to reverse the hill effect. Here we use the simplest linear monotonically increasing assumption to construct the loss function:

$$\mathcal{L}_{LA} = \begin{cases} \|U[n_c] - V_{th} \frac{n_c}{n^*}\|_2^2, & \text{if condition}_n; \\ 0, & \text{otherwise,} \end{cases} \quad (9)$$

where n^* denotes the expected spike timestep and n_c denotes the current spike timestep, condition_n corresponds to the nonlinear monotonically increasing condition that satisfies: $U[n_c] \geq$

Algorithm 1 Adaptive Event Slicing Process

Input: SNN model, input event \mathcal{E} , the number of event cell N , the previous spike index n_p , list \mathcal{D}_{rep} .
Initialize: $n_p = 0$.
for all $n = 0, 1, 2, \dots, N - 1$ **do**
 $S_{out} = \text{SNN}(C[n])$.
 if $S_{out} == 1$ **then**
 $n_c = n$.
 $T_{group} = \bigcup_{n=n_p+1}^{n_c} f_{time}(n)$.
 $\mathcal{E}_{group} = G_{\pm}(x, y, t, T_{group})$.
 $\mathcal{D}_{n_c} = \mathcal{F}(\mathcal{E}_{group})$.
 Append \mathcal{D}_{n_c} to \mathcal{D}_{rep} .
 $n_p = n_c + 1$.
 end if
end for
Return: \mathcal{D}_{rep} .

Algorithm 2 Feedback-Update Training Strategy

Input: SNN model, pretrained ANN model, ANN training dataset \mathcal{D}_A , SNN training dataset \mathcal{D}_S , total epoch E_{train} , epoch to start finetuning e_f , alpha α , learning rate η .
for all $e = 1, 2, \dots, E_{train}$ epoch **do**
 for all event batch $d_i = d_1, d_2, \dots, d_{N_s}$ in \mathcal{D}_S **do**
 Feed d_i into SNN until it spikes at time step n_c .
 Generate event candidates to ANN to get feedback.
 Calculate loss function $\mathcal{L}_{SPA} = \mathcal{L}_{Mem} + \mathcal{L}_{LA}$.
 Backpropagate and update SNN parameters.
 end for
 $\alpha \leftarrow \alpha - 2 \cdot \eta \sum_i^{N_s} (n^{*i} - n_c^i) / N_s$.
 if $e > e_f$ **then**
 Split \mathcal{D}_A into \mathcal{D}'_A by adaptive event slicing process.
 Finetune ANN parameters on \mathcal{D}'_A .
 end if
end for

$U[n^*]$ & $n_c < n^*$. We expect the membrane potential at n_s to reach $\frac{n_c}{n^*} V_{th}$ for the latter membrane potential at the n^* to reach V_{th} in a linearly increasing form. More explanations are provided in Appendix I.2.

Combined with Mem-Loss and LA-Loss, we defined the SPA-Loss, which guides the adaptive event slicing process in subsequent experiments: $\mathcal{L}_{SPA} = \mathcal{L}_{Mem} + \mathcal{L}_{LA}$.

3.3.3 Dynamic Hyperparameter Tuning

Although controlling the SNN to spike at a desired location can be achieved through the combination of Mem-Loss and LA-Loss, the utilization of varying α values (Eq. 8) may result in significant fluctuations in experimental results. We have the following observation.

Observation 2. The larger the α , the earlier the SNN tends to fire spikes; and vice versa.

A larger α in Eq. 8 implies a higher pre-momentary membrane potential, which results in an earlier spike. Taking the larger- α scenario in Fig. 3(c) as an example, if the SNN is expected to activate a spike at a later timestep, the larger α prevents the actual spike from being delayed. Thus, we need to decrease α , causing the expected spike time to shift earlier. This concludes that the update direction of the hyperparameter α should be consistent with the update direction of the desired spiking index.

Observation 3. A fixed α leads to significant variations in performance across different tasks.

As illustrated in Fig. 3(b) and (c), the same α varies significantly on different downstream models, which makes it difficult to set the hyperparameter alpha in advance.

To address the above issues, we design a learning-based hyperparameter tuning method for updating α (in Alg. 2). More details are provided in Appendix I.3.

3.4 Feedback-Update Strategy through SNN-ANN Cooperation

Based on the methods proposed in the previous sections: if the desired trigger time n^* is given, the SNN is able to accurately accomplish the event slicing under the guidance of the SPA-Loss function. In this section, we focus on how to obtain the desired spike time n^* through the downstream ANN feedback. We thus propose a feedback-update strategy that enables the SNN to slice events when the downstream ANN model achieves optimal performance. By receiving real-time feedback from the downstream model and updating n^* , this strategy ultimately enhances task performance.

Particularly, when SNN processes the input event and triggers a spike at time n_c , it returns a spike output sequence $S = [0, \dots, 0, 1, 0, \dots]$, where 1 is at n_c -th. We first perform a *neighborhood search* to obtain $2d + 1$ candidate event representation with the index in $\{n_c - d, \dots, n_c + d\}$: $\{D_{n_c-d}, \dots, D_{n_c+d}\}$, where $D_{n_c+i} = \mathcal{F}(G_{\pm}(x, y, t, \{t \in [t_0 + (n_p + 1)\delta t, t_0 + (n_c + 1 + i)\delta t]\}))$. We then choose a downstream model \mathcal{M} (e.g., object tracker or recognizer) and feed the candidate event representations into it to obtain feedback y :

$$y = \mathcal{L}_{\mathcal{M}}(C[n_c - d]) \oplus \dots \oplus \mathcal{L}_{\mathcal{M}}(C[n_c + d]), \quad (10)$$

where $\mathcal{L}_{\mathcal{M}}(\cdot)$ returns the output loss of \mathcal{M} and \oplus concatenates these losses into $y \in \mathbb{R}^{2d+1}$. We choose the model loss as the feedback since it directly reflects the quality of inputs. We can then generate the desired spike index n^* by: $n^* = \arg \min_i (y[i])$, where $\arg \min$ extracts the index with




Figure 4: (a) Experiments on comparing different loss functions on a simple event slicing task. Our proposed Mem-Loss and LA-Loss require only a small number of iterations to supervise the SNN to activate spikes at the desired time steps; (b) Experiments on different hyperparameter settings. Our dynamic tuning method can stably converge towards the optimal spiking time (colored in green). In contrast, using a fixed α results in unstable training and challenges in finding the optimal point.

the best (minimal loss) feedback, which in turn guides the dynamic slicing process using SPA-Loss. After training the SNN, the ANN is then updated by feeding the newly-sliced events, thus forming an SNN-ANN cooperation process.

Feedback-Update Strategy. This strategy employs a two-stage iterative approach. In the first stage, the ANN offers real-time feedback to the adaptive slicing process for training SNN. In the second stage, the trained SNN provides a newly-sliced event to finetune the ANN. The process then iterates back to the first stage. This strategy provides a novel SNN-ANN cooperation paradigm which establishes a strong connection between raw data and the downstream model. We summarize the feedback-update strategy in Alg. 2.

4 Experiments

To evaluate the effectiveness of our proposed method, we set up two-level experiments. In the beginner’s arena, we expect the SNN to find the exact slicing time with the simulated event inputs. In the expert’s arena, we conduct experiments on event-based object tracking and image recognition. Details of experiment settings and more experimental analyses are presented in the Appendix. Our code will be available at <https://github.com/AndyCao1125/SpikeSlicer>.

4.1 Beginner’s Arena: Event Slicing in Simple Tasks

We first conduct some entry-level tasks to validate the effectiveness of SPA-Loss. We set up the task: *Input N randomized event cells, expect the SNN to slice at a specified time step n^* and there exists a certain probability of interfering with SNN to slice at other time steps.*

We compare the SPA-Loss function with common MSE-Loss and CE-Loss. The experiment setting is detailed in Appendix K. As depicted in Fig 4(b), the SPA-Loss successfully supervises the SNN to activate spikes at the desired time steps. In particular, SPA-Loss requires only a few iterations (<400) to supervise the SNN to fire spikes correctly. In contrast, MSE-Loss can only succeed at certain time steps, and CE-Loss cannot even accomplish the task. In addition, using both Mem-Loss and LA-Loss yields smoother results compared to using Mem-Loss alone. To summarize, the beginner’s arena preliminarily tests the effectiveness of the SPA-Loss and paves the way for subsequent experiments.

4.2 Expert’s Arena: Mastering Adaptive Event Slicing with SNN-ANN Collaboration

After a successful challenge in the beginner’s arena, we move on to the expert arena. Here we use a low-energy SNN to adaptively process the event data on complex downstream tasks:

Event-based Object Tracking. Since the tracking task is highly sensitive to temporal information, dynamic event slicing is of great importance. Tab. 1 shows that the tracking performances under the SpikeSlicer have a significant improvement in terms of representative success rate (RSR), representative precision rate (RPR), and overlap precision (OP). For instance, TransT’s performance on $OP_{0.50}$ and $OP_{0.75}$ improved by 9.8% and 27.7% compared to its original result under the HDR scenario. When compared to results on the fixed event (the number of the fixed-sliced event is aligned with the number of dynamic-sliced events to ensure a fair comparison), our method achieves favorable gains in the overall RSR, *i.e.*, 62.4 vs. 51.0.

Event-based Recognition. We also conduct experiments in event-based recognition to evaluate the effectiveness of our proposed method. As depicted in Tab. 2, our method has a significant

Table 1: Quantitative comparison on FE108. There are four challenging scenarios, including high dynamic range (HDR), low light (LL), fast motion with and without motion blur (FWB & FNB) and all testing datasets (ALL). $[method]+SpikeSlicer$ represents the results based on our adaptive event slicing method and the results of $[method]$ are tested on the original fixed-sliced event dataset from [3]. To ensure a fair comparison, *fix event* indicates that the model is tested on a dataset of fixed-sliced event frames, where the number of fixed event frames is the same as the number of dynamically sliced event frames by using SpikeSlicer. Best performances are denoted by **deep green**.

Methods	HDR				LL				FWB				FNB				ALL			
	RSR	OP _{.50}	OP _{.75}	RPR	RSR	OP _{.50}	OP _{.75}	RPR	RSR	OP _{.50}	OP _{.75}	RPR	RSR	OP _{.50}	OP _{.75}	RPR	RSR	OP _{.50}	OP _{.75}	RPR
SiamFC++ [29]	15.3	15.0	1.3	25.2	13.4	8.7	0.8	15.3	28.6	36.3	6.0	48.2	36.8	42.7	7.4	63.1	23.8	26.0	3.9	39.1
KYS [30]	15.7	14.5	5.2	23.0	12.0	8.0	1.1	18.0	47.0	63.9	14.8	73.3	36.9	44.5	15.2	57.9	26.6	30.6	9.2	41.0
CLNet [31]	30.0	33.5	9.6	48.3	13.7	6.0	0.9	23.6	52.9	71.2	23.3	80.3	40.8	46.3	14.2	67.7	34.4	39.1	11.8	55.5
DiMP [28]	41.8	50.0	17.9	62.7	45.6	52.8	11.2	69.5	69.4	94.7	37.1	99.7	60.5	75.6	29.3	93.2	52.6	65.4	23.4	79.1
DiMP (<i>fixed event</i>)	53.3	68.2	21.4	81.6	67.6	86.3	43.1	95.0	49.7	45.4	5.88	80.5	49.6	59.4	23.7	75.3	53.8	64.3	19.5	82.4
DiMP+SpikeSlicer	52.3	65.5	21.4	78.7	69.4	92.4	43.2	97.5	52.4	64.8	26.6	78.8	62.2	81.0	21.8	97.2	56.0	71.4	23.3	85.1
PrDiMP [32]	44.3	52.8	19.6	66.3	44.6	48.2	8.9	69.5	67.0	89.9	33.6	99.7	60.6	75.8	29.7	93.3	53.0	65.0	23.3	80.5
PrDiMP (<i>fixed event</i>)	41.2	49.8	18.4	66.1	42.7	45.2	12.5	87.1	62.4	85.8	21.3	90.4	47.6	58.3	18.5	77.0	48.0	61.2	19.2	78.6
PrDiMP+SpikeSlicer	53.3	67.2	22.1	79.7	70.1	94.5	41.8	97.4	70.7	88.3	55.5	94.3	56.3	70.5	25.9	83.9	59.2	75.1	29.4	86.7
TransT [27]	54.0	68.5	22.0	81.7	67.9	88.4	40.4	96.7	67.0	92.4	30.5	99.6	58.0	73.6	27.5	87.2	58.1	74.7	25.7	87.7
TransT (<i>fixed event</i>)	51.4	67.8	11.1	81.2	63.2	80.2	28.3	89.3	41.5	28.0	2.50	57.7	50.6	57.9	12.7	78.9	51.0	59.0	12.0	78.8
TransT+SpikeSlicer	57.7	75.2	28.1	82.6	70.7	93.6	42.7	99.0	74.9	97.7	61.1	98.6	58.7	75.6	29.6	84.6	62.4	81.2	34.1	88.9

Table 2: Quantitative comparison on DVS-Gesture, N-Caltech101, DVS-CIFAR10 and SL-Animals. *Random* and *Fix* denote that the input events are randomly sliced and fixed sliced, respectively. Instead, our method slices the event stream adaptively.

Method	DVS-Gesture [33]			N-Caltech101 [34]			DVS-CIFAR10 [35]			SL-Animals [36]		
	Random	Fix	Ours	Random	Fix	Ours	Random	Fix	Ours	Random	Fix	Ours
ResNet-18	93.06	93.40	94.79	77.80	73.37	79.86	78.91	77.73	81.15	85.71	83.93	88.39
ResNet-34	95.14	93.40	96.18	78.77	76.08	82.54	80.57	79.39	82.23	86.61	87.50	89.93
Swin-S	88.19	89.93	91.67	86.30	81.76	87.30	81.54	79.79	83.01	74.11	56.25	75.45

improvement over the fixed-sliced method, with an accuracy improvement of 2.78% and 6.46% by using ResNet-34 in DVS-Gesture and N-Caltech101, respectively. To verify that the results of our adaptive slicing method are not biased due to randomness, we add the random-slice baselines for comparison, in which the event stream is randomly sliced into event groups and fed into the ANN for training. Our method also yields better performance compared with the random-slice results.

Visualization of Adaptive Event Slicing. We visualize the tracking results to demonstrate that the dynamic slicing method is able to adapt to various motion scenarios. As shown in Fig. 5, our method obtains better tracking performance compared to fixed event inputs, *i.e.*, the position of the prediction box is more accurate. Additionally, our dynamic event slicing method can achieve (1) *edge enhancement* and (2) *redundancy removal* to refine the event data under different tracking scenarios. However, the fixed-slice approach adopts the same slicing strategy for each event stream, leading to performance degradation.

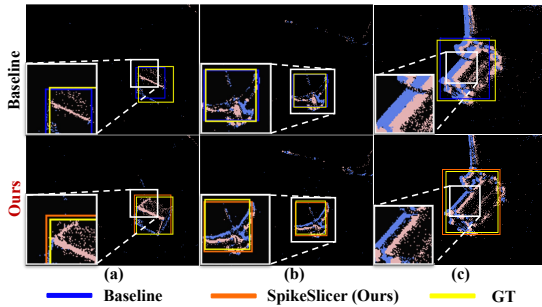


Figure 5: Visualization results on FE108 dataset. The white box denotes the zoom-in area. Our adaptive event slicing method provides better tracking performance than fixed counterparts while enabling edge enhancement (a,b) and redundancy removal (c).

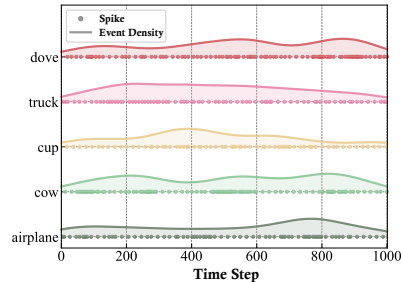


Figure 6: Visualization of spike splitting points with corresponding event density. The results of our dynamic splitting are **matched** to real event information. The higher the event density, the SNN splits more frequently; conversely, the SNN splits more sparsely.

Table 3: Comparison of Efficiency and Speed. The comparison includes the number of operations (OPs) and tracking speed per image without image processing time.

Models	OPs (G)	Energy (mJ)	Speed (s / img)	Performance
ANN w/o SpikeSlicer	56.36	259.26	0.045	51.0
ANN w/ SpikeSlicer	57.09	260.11	0.060	62.4

Table 4: Experiments on different event representations with fixed or dynamic slicing methods. Our method yields significant improvement when using different event representation methods, proving SpikeSlicer’s effectiveness as a plug-and-play event slicer.

FE108	Time Surface		Event Spike Tensor		Voxel Grid	
Slice Method	RSR	RPR	RSR	RPR	RSR	RPR
Fix Slice	57.5	85.7	50.4	81.6	51.0	78.8
SpikeSlicer (ours)	59.5	86.8	54.5	85.6	62.4	88.9

DVS-Gesture	Event Frame	Event Spike Tensor	Voxel Grid
Slice Method	Accuracy	Accuracy	Accuracy
Fix Slice	93.75%	93.75%	88.54%
SpikeSlicer (ours)	94.79%	95.49%	89.24%

Table 5: Ablation studies for evaluating the proposed loss function on the SL-Animals dataset.

Slice Method	\mathcal{L}_{Mem}	\mathcal{L}_{LA}	ResNet-18	ResNet-34
Fix Slice	X	X	83.93	87.50
SpikeSlicer	✓	X	87.50 (+3.57)	88.52 (+1.02)
SpikeSlicer	✓	✓	88.39 (+4.46)	89.73 (+2.23)

Table 6: Experiments with different event cell numbers N . The resulting sliced event group always has a similar time interval in various N conditions.

N	15	20	25
Avg Spike Time	2.42	3.15	4.77
Time Duration	16.13%	15.75%	19.08%

4.3 Analysis of the Adaptive Slicing Method

Analysis of Spike Splitting Time and Event Density. To evaluate the effectiveness of our proposed method, we conducted a detailed visualization analysis, as depicted in Fig. 6, examining the relationship between the locations of split points and the corresponding event stream densities. The definition of event density is detailed in the Appendix B. The analysis reveals a clear **match** between the positions of cuts made by SNN and the respective event density. Specifically, the SNN tends to perform more frequent cuts in regions of higher event density, while conversely, regions with lower event density experienced fewer cuts. These findings indicate that the dynamic cutting method is effectively adaptive to the varying information density within the event stream.

Analysis of Efficiency and Speed. We evaluate the performance of the ANN tracker (TransT) with or without the SpikeSlicer in terms of energy consumption and processing speed. Detailed calculations for energy consumption are in Appendix M. In Tab. 3, the SpikeSlicer stands out as a plug-and-play option. The introduction of the SpikeSlicer incurs a marginal energy increase of 0.85 mJ and a small reduction in processing speed. However, these costs are offset by a significant 22.3% improvement in the RSR metric.

Evaluation of the Dynamic Hyperparameter Tuning. We use different α settings in tracking experiments to study the effect of hyperparameters. Fig. 4(b) shows that the fixed α setting does not allow the spiking time step to reach the desired time step accurately, while our dynamic α tuning allows SNN to spike at the desired time step, leading to a better event-slicing process.

4.4 Ablation Study

Ablation Study on Event Representation Method. Since SpikeSlicer works as a plug-and-play event slicing method suitable for multiple event representations, we thus conduct our method through different event representation methods to further evaluate the effectiveness. Tab. 4 demonstrates that, across various forms of event representation, the dynamic splitting method yields performance enhancements in both tracking and recognition tasks when compared to the fixed splitting method.

Ablation Study on Different Components of SPA-Loss. The results in Tab. 5 reveal that the adaptive slicing method achieves accuracy improvement, and all of our proposed loss functions (including \mathcal{L}_{Mem} and \mathcal{L}_{LA}) contribute to the performance, showing the effectiveness of our method.

Ablation Study on Event Cell Number. Considering that the number of event cells N may affect the SNN’s decision on event slicing, we examine the stability of the sliced event group by varying the size of N . Tab. 6 shows that the average spike time of the SNN varies for different N , but the time duration (*i.e.*, Avg Spike/ N) of the resulting event groups is stable. This verifies that the SNN can effectively make cuts based on event information rather than making decisions based on the number of inputs alone. More details are provided in Appendix P.

5 Limitation

We summarize the limitations of this work as follows: Firstly, the SpikeSlicer process involves multi-stage SNN-ANN training, which leaves substantial room for improvement in the adaptive slicing strategy. Secondly, for recognition tasks, we convert the event stream into a single-frame representation to obtain accurate supervisory signals. This approach could be refined in the future to enable SpikeSlicer to slice stream events into multi-frame representations, which are the mainstream format. Thirdly, our experiments are conducted on GPUs; however, the most suitable hardware for SNNs would be brain-inspired chips. In addition, dynamic events need to be generated and processed in real-time during inference, rather than fixed generation in advance. As a result, conducting experiments on GPUs may lead to slower overall inference speeds. Extending this paradigm to brain-inspired chips with asynchronous event input is an interesting direction worth exploring in the future.

6 Conclusion

In this work, we proposed SpikeSlicer, a novel event processing method that splits event streams adaptively. SpikeSlicer utilizes a spiking neural network (SNN) as an event trigger, which determines the slicing time according to the generated spikes. To achieve accurate slicing, we designed the Spiking Position-aware Loss (SPA-Loss) which guides the SNN to trigger spikes at the desired time step. In addition, we proposed a Feedback-Update training strategy that allows the SNN to make accurate slicing decisions based on the ANN feedback. Extensive experiments have demonstrated the effectiveness of SpikeSlicer in yielding performance improvement in event-based object tracking and recognition tasks. In the future, we will assess SpikeSlicer’s suitability for other event-based tasks, and devise more efficient training strategies for the SNN-ANN cooperative framework to optimize real-time processing in the future.

References

- [1] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conradt, Kostas Daniilidis, et al. Event-based vision: A survey. *IEEE TPAMI*, 44(1):154–180, 2020.
- [2] Jiqing Zhang, Bo Dong, Haiwei Zhang, Jianchuan Ding, Felix Heide, Baocai Yin, and Xin Yang. Spiking transformers for event-based single object tracking. In *CVPR*, pages 8801–8810, 2022.
- [3] Jiqing Zhang, Xin Yang, Yingkai Fu, Xiaopeng Wei, Baocai Yin, and Bo Dong. Object tracking by jointly exploiting frame and event domain. In *ICCV*, pages 13043–13052, 2021.
- [4] Jiyuan Zhang, Lulu Tang, Zhaofei Yu, Jiwen Lu, and Tiejun Huang. Spike transformer: Monocular depth estimation for spiking camera. In *ECCV*, pages 34–52. Springer, 2022.
- [5] Yeongwoo Nam, Mohammad Mostafavi, Kuk-Jin Yoon, and Jonghyun Choi. Stereo depth from events cameras: Concentrate and focus on the future. In *CVPR*, pages 6114–6123, 2022.
- [6] Amos Sironi, Manuele Brambilla, Nicolas Bourdis, Xavier Lagorce, and Ryad Benosman. Hats: Histograms of averaged time surfaces for robust event-based object classification. In *CVPR*, pages 1731–1740, 2018.
- [7] Raymond Baldwin, Ruixu Liu, Mohammed Mutlaq Almatrafi, Vijayan K Asari, and Keigo Hirakawa. Time-ordered recent event (tore) volumes for event cameras. *IEEE TPAMI*, 2022.
- [8] Lin Wang, Yo-Sung Ho, Kuk-Jin Yoon, et al. Event-based high dynamic range image and very high frame rate video generation using conditional generative adversarial networks. In *CVPR*, pages 10081–10090, 2019.
- [9] Daniel Gehrig, Antonio Loquercio, Konstantinos G Derpanis, and Davide Scaramuzza. End-to-end learning of representations for asynchronous event-based data. In *ICCV*, pages 5633–5643, 2019.
- [10] Ana I Maqueda, Antonio Loquercio, Guillermo Gallego, Narciso García, and Davide Scaramuzza. Event-based vision meets deep learning on steering prediction for self-driving cars. In *CVPR*, pages 5419–5427, 2018.
- [11] Lin Zhu, Xiao Wang, Yi Chang, Jianing Li, Tiejun Huang, and Yonghong Tian. Event-based video reconstruction via potential-assisted spiking neural network. In *CVPR*, pages 3594–3604, 2022.

- [12] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Unsupervised event-based learning of optical flow, depth, and egomotion. In *CVPR*, pages 989–997, 2019.
- [13] Yansong Peng, Yueyi Zhang, Peilin Xiao, Xiaoyan Sun, and Feng Wu. Better and faster: adaptive event conversion for event-based object detection. In *AAAI*, volume 37, pages 2056–2064, 2023.
- [14] Jianing Li, Jia Li, Lin Zhu, Xijie Xiang, Tiejun Huang, and Yonghong Tian. Asynchronous spatio-temporal memory network for continuous event-based object detection. *IEEE TIP*, 31:2975–2987, 2022.
- [15] Jiqing Zhang, Yuanchen Wang, Wenxi Liu, Meng Li, Jinpeng Bai, Baocai Yin, and Xin Yang. Frame-event alignment and fusion network for high frame rate tracking. In *CVPR*, pages 9781–9790, 2023.
- [16] Inwoo Hwang, Junho Kim, and Young Min Kim. Ev-nerf: Event based neural radiance field. In *WACV*, pages 837–847, 2023.
- [17] Viktor Rudnev, Mohamed Elgharib, Christian Theobalt, and Vladislav Golyanik. Eventnerf: Neural radiance fields from a single colour event camera. In *CVPR*, pages 4992–5002, 2023.
- [18] Ziqing Wang, Yuetong Fang, Jiahang Cao, Qiang Zhang, Zhongrui Wang, and Renjing Xu. Masked spiking transformer. In *ICCV*, pages 1761–1771, October 2023.
- [19] Jiahang Cao, Ziqing Wang, Hanzhong Guo, Hao Cheng, Qiang Zhang, and Renjing Xu. Spiking denoising diffusion probabilistic models. In *WACV*, pages 4912–4921, 2024.
- [20] Eric Hunsberger and Chris Eliasmith. Spiking deep networks with lif neurons. *arXiv preprint arXiv:1510.08829*, 2015.
- [21] Anthony N Burkitt. A review of the integrate-and-fire neuron model: I. homogeneous synaptic input. *Biological Cybernetics*, 95:1–19, 2006.
- [22] Patrick Bardow, Andrew J Davison, and Stefan Leutenegger. Simultaneous optical flow and intensity estimation from an event camera. In *CVPR*, pages 884–892, 2016.
- [23] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham China, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1):82–99, 2018.
- [24] Filipp Akopyan, Jun Sawada, Andrew Cassidy, Rodrigo Alvarez-Icaza, John Arthur, Paul Merolla, Nabil Imam, Yutaka Nakamura, Pallab Datta, Gi-Joon Nam, et al. Truenorth: Design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *IEEE TCAD*, 34(10):1537–1557, 2015.
- [25] Qinyi Wang, Yexin Zhang, Junsong Yuan, and Yilong Lu. Space-time event clouds for gesture recognition: From rgb cameras to event cameras. In *WACV*, pages 1826–1835, 2019.
- [26] Zhiyu Zhu, Junhui Hou, and Xianqiang Lyu. Learning graph-embedded key-event back-tracing for object tracking in event clouds. *NeurIPS*, 35:7462–7476, 2022.
- [27] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *CVPR*, 2021.
- [28] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *ICCV*, 2019.
- [29] Yinda Xu, Zeyu Wang, Zuoxin Li, Ye Yuan, and Gang Yu. Siamfc++: Towards robust and accurate visual tracking with target estimation guidelines. In *AAAI*, 2020.
- [30] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Know your surroundings: Exploiting scene information for object tracking. In *ECCV*, 2020.
- [31] Xingping Dong, Jianbing Shen, Ling Shao, and Fatih Porikli. Clnet: A compact latent network for fast adjusting siamese trackers. In *ECCV*, 2020.
- [32] Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *CVPR*, 2020.
- [33] Arnon Amir, Brian Taba, David Berg, Timothy Melano, Jeffrey McKinstry, Carmelo Di Nolfo, Tapan Nayak, Alexander Andreopoulos, Guillaume Garreau, Marcela Mendoza, et al. A low power, fully event-based gesture recognition system. In *CVPR*, pages 7243–7252, 2017.

- [34] 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.
- [35] Hongmin Li, Hanchao Liu, Xiangyang Ji, Guoqi Li, and Luping Shi. Cifar10-dvs: an event-stream dataset for object classification. *Frontiers in Neuroscience*, 11:244131, 2017.
- [36] Ajay Vasudevan, Pablo Negri, Camila Di Ielsi, Bernabe Linares-Barranco, and Teresa Serrano-Gotarredona. SI-animals-dvs: event-driven sign language animals dataset. *Pattern Analysis and Applications*, pages 1–16, 2022.
- [37] Arjun Roy, Manish Nagaraj, Chamika Mihiranga Liyanagedera, and Kaushik Roy. Live demonstration: Real-time event-based speed detection using spiking neural networks. In *CVPR*, pages 4080–4081, 2023.
- [38] Alberto Viale, Alberto Marchisio, Maurizio Martina, Guido Maserà, and Muhammad Shafique. Carsnn: An efficient spiking neural network for event-based autonomous cars on the loihi neuromorphic research processor. In *IJCNN*, pages 1–10. IEEE, 2021.
- [39] Fangwen Yu, Yujie Wu, Songchen Ma, Mingkun Xu, Hongyi Li, Huanyu Qu, Chenhang Song, Taoyi Wang, Rong Zhao, and Luping Shi. Brain-inspired multimodal hybrid neural network for robot place recognition. *Science Robotics*, 8(78):eabm6996, 2023.
- [40] Jesse Hagenaaars, Federico Paredes-Vallés, and Guido De Croon. Self-supervised learning of event-based optical flow with spiking neural networks. *NeurIPS*, 34:7167–7179, 2021.
- [41] Man Yao, Huanhuan Gao, Guangshe Zhao, Dingheng Wang, Yihan Lin, Zhaoxu Yang, and Guoqi Li. Temporal-wise attention spiking neural networks for event streams classification. In *ICCV*, pages 10221–10230, 2021.
- [42] Xavier Lagorce, Garrick Orchard, Francesco Galluppi, Bertram E Shi, and Ryad B Benosman. Hots: a hierarchy of event-based time-surfaces for pattern recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(7):1346–1359, 2016.
- [43] Priyadarshini Panda, Sai Aparna Aketi, and Kaushik Roy. Toward scalable, efficient, and accurate deep spiking neural networks with backward residual connections, stochastic softmax, and hybridization. *Frontiers in Neuroscience*, 14:653, 2020.
- [44] Man Yao, Guangshe Zhao, Hengyu Zhang, Yifan Hu, Lei Deng, Yonghong Tian, Bo Xu, and Guoqi Li. Attention spiking neural networks. *IEEE TPAMI*, 2023.
- [45] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021.
- [46] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020.
- [47] Junho Kim, Jaehyeok Bae, Gangin Park, Dongsu Zhang, and Young Min Kim. N-imagenet: Towards robust, fine-grained object recognition with event cameras. In *ICCV*, pages 2146–2156, 2021.

Appendix

A Details of our Motivations

To clarify the motivation behind our dynamic event stream slicing algorithm, this section tells the details.

A.1 Motivation for Proposing a Dynamic Event Stream Slicing Algorithm

The process of event-to-representation conversion is mainly divided into two steps: **Step 1. Slice the raw event stream into multiple sub-event stream**, and **Step 2. convert these sub-event streams into representations using various event representation methods**. While much work has focused on optimizing event representation (Step 2) to extract better event information, including time surface and EST, they do not address the issues arised with fixed slicing (*e.g.*, resulting non-uniform event in scenarios with changing motion speed). Despite event slicing being a small part of the overall pipeline, it is a critical point. This is because the event stream is very sensitive to slicing, and the model performance fluctuates very much for different slicing methods, as proved by extensive experiments in Appendix C.

To better address this issue, we introduced the dynamic slicing method SpikeSlicer. Meanwhile, SpikeSlicer is guided by downstream task feedback to ensure that the new sub-streams could enhance downstream task performance.

A.2 Motivation for for Using SNN as a Slicing Trigger

The reason why we choose SNN as the event slicing trigger is twofold:

- Utilizing SNNs on neuromorphic hardware for processing event streams is low-energy and low-latency [23, 24].
- Deployed on neuromorphic hardware, SNNs can process event streams asynchronously [37–39], conserving energy when there is no data input—a capability that GPUs, operating synchronously, lack.

Due to the aforementioned reasons, there is a considerable amount of research [40, 41, 11] employing Spiking Neural Networks (SNNs) for event data. Although these SNNs are simulated on GPU platforms, the models resulting from such simulations could be deployed on neuromorphic hardware.

A.3 Contribution for Using SNN as a Slicing Trigger

We propose a new cooperative paradigm where SNN acts as an efficient, low-energy data processor to assist the ANN in improving downstream performance. This is a brand-new SNN-ANN cooperation way, paving the way for future event-related implementation on neuromorphic chips.

B Definition of Event Density

In our experiments, we investigated the relationship between the location of the split point determined by the SNN and the density of the corresponding event stream. The event density, denoted as $D(t)$, measures the concentration of events in a given event stream over time. It is mathematically defined as the rate at which events occur per timestep, expressed as a function of:

$$D(t) = \frac{\delta N}{\delta t} \quad (11)$$

where $D(t)$ is the event density at timestep t , δN is the number of events occurring within a small time interval δt around t . This definition enables a precise quantification of event concentration, offering insights into the temporal distribution of events at any given moment. Our empirical analysis shows that there is a significant correlation between the split points of the SNN and the event density. Notably, in regions of higher event density, the SNN exhibits a tendency to perform more frequent split. In contrast, in regions with lower event density, the number of split is lower. This behavior emphasizes the adaptability of our proposed dynamic slicing method, SpikeSlicer, to the fluctuating information density in the event stream.

C Sensitivity Analysis of Fixed Event Slicing Method

To demonstrate that events are sensitive to slicing by fixed methods, and to emphasize the importance of proposing a dynamic event slicing approach, we have conducted a total of 60 experiments with different models to investigate the impact of different slicing techniques and different numbers of slices on the performance in downstream tasks.

In our experiment, we employed two fixed slicing methods: (1). *Slicing with a fixed number of events*, and (2). *Slicing with a fixed duration*. N denotes the number of resulting event slices. Experimental results are detailed as follows in Tab. 7 and Fig. 7.

Table 7: The sensitivity analysis of fixed event slicing on N-Caltech101. The results demonstrate that the event is sensitive to the fixed slicing method (slicing by fixed time or event count), thereby affirming the need for proposing a dynamic slicing method.

N-Caltech101	N	2	4	6	8	10	12	14	16	18	20	22	24	26	28	30	Mean	Var
ResNet18	Fixed Count	70.96	75.26	75.39	75.30	76.09	73.95	74.09	73.80	76.40	75.39	75.45	73.60	71.94	71.01	71.17	73.98	3.33
ResNet18	Fixed Time	62.90	72.64	76.38	74.48	74.91	73.70	74.30	74.69	76.95	74.75	74.46	74.42	71.61	71.52	69.69	73.16	10.80
ResNet34	Fixed Count	72.19	75.55	76.98	78.22	77.14	77.40	76.78	76.90	78.14	77.06	76.91	74.85	74.76	76.91	73.07	76.19	2.90
ResNet34	Fixed Time	65.42	75.92	78.29	78.20	78.48	76.22	77.76	76.57	75.94	76.80	76.61	75.91	74.76	74.19		75.74	9.15

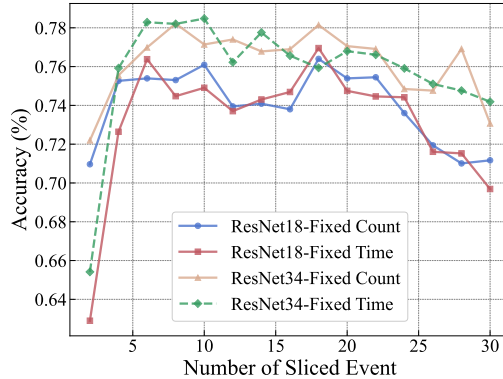


Figure 7: Visualization of sensitivity analysis on N-Caltech101 dataset. The fluctuations in accuracy for different numbers of sliced event with different fixed slicing methods are significant, demonstrating that events are very sensitive to fixed slicing methods.

The results indicate significant fluctuations (large variance) in downstream performance based on the slicing method and the number of slices used. We believe this addition effectively demonstrates the sensitivity of event streams to slicing techniques, **confirming the need for our motivation to propose dynamic slicing of event streams**. Additionally, the accuracy achieved using the dynamic slicing method (82.54% by ResNet34) surpasses that of any fixed slicing approach (with the highest being 78.48%), further substantiating the efficacy of the dynamic method in our study.

D Illustration of SpikeSlicer vs. Fixed Slicing Method

In order to more intuitively show the difference between our dynamic event slicing method and the traditional fixed slicing method, we specifically illustrate these methods through Fig. 1.

Fig. 1(a) denotes that each resulting sliced event has the same duration, and Fig. 1(b) denotes that the number of event points contained in each resulting sliced event is the same. Since event stream are usually unevenly distributed, a fixed cutting method often leads to non-uniform event information (*e.g.*, in scenarios with changing motion speed). In contrast, our approach decides the optimal slice position through feedback from downstream ANN by using an SNN as the event slicing trigger.

E Difference between Event Slicing and Event Representation

It is worth noting that **our work focuses on the slicing of the event stream rather than focusing on event representation**. Event representation refers to the process of event information extraction

Table 8: Experiments on different event representations with fixed (including fixed time and fixed event count) or dynamic slicing methods. Our SpikeSlicer yields significant improvement when using different event representation methods.

DVSGesture	Event Frame	Event Spike Tensor	Voxel Grid
Fix Duration	93.75%	93.75%	88.54%
Fix Event Count	93.06%	94.79%	88.19%
SpikeSlicer (ours)	94.79%	95.49%	89.24%

that is performed after the event stream has been sliced into sub-event stream, and the resulting event representation meets the neural network input requirements. Thus, our dynamic slicing process and event representation can be used at the same time, either better slicing or representation method benefits the feature extraction with neural network, thus improving performance.

To validate the effectiveness of our slicing approach, we supplement the event-based recognition task below. We compare the downstream performance of three different event representation methods (including Event Frame [10], Event Spike Tensor (EST [9]) and Voxel Grid [12]) on the DVSGesture dataset under fixed slicing and dynamic slicing method in Tab. 8.

F Definition of Raw Event Group

Based on the definition of event field from [9], we here define a general version of raw event group representation as a mapping $\mathcal{M} : \mathcal{E} \mapsto \mathcal{T}$ between the set \mathcal{E} and a tensor \mathcal{T} :

Definition 1 (Raw event group). *Based on a measurement $\mathbb{1}_{\text{condition}}$, raw event group are grid-like tensors defined in continuous space and time:*

$$G_{\pm}(x, y, t, \text{condition}) = \sum_{e_k \in \mathcal{E}_{\pm}} \mathbb{1}_{\text{condition}}(x, y, t) \delta(x - x_k, y - y_k) \delta(t - t_k), \quad (12)$$

where \pm denotes the event polarity; $\mathbb{1}_{\text{condition}}$ sets the specific approach of representation to each event, e.g., $\text{condition} = \{\sum_{e_k} = M\}$ denotes the number of event points in each grid-like tensor is fixed to M , i.e., slicing event stream \mathcal{E} by *event count*; or $\text{condition} = \{\Delta t = t_x\}$ denotes the time interval of event points in each grid-like tensor is fixed to t_x . G_{\pm} are grid tensors with $x \in \{0, 1, \dots, W - 1\}$, $y \in \{0, 1, \dots, H - 1\}$, and $t \in \{t_0, t_0 + \Delta t_{x1}, \dots, t_0 + B\Delta t_{xB}\}$, where t_0 is the first time stamp, Δt_{xi} is the bin size determined by the splitting condition, and B is the number of temporal bins. Eq. (12) converts raw event into grid-like raw event group by a Dirac pulse [9] in the space-time manifold. The resulting G_{\pm} gives a continuous time representation of \mathcal{E} which preserves the event’s information.

However, if such raw event groups are then converted into event representation (e.g. voxel grid [22], time surface [42]), the generated event representation is imprecise due to the fact that the process of $\mathbb{1}_{\text{condition}}$ is fixed, leading to both spatial and temporal information loss. The main objective of this study is to solve the problem of fixed slicing of the event stream and to provide a dynamic segmentation scheme.

G Reason for Using No-reset Membrane Potential

We first recall the definition of original membrane potential $V[n]$:

$$V[n] = \beta V[n - 1] + \gamma I[n], \quad (13)$$

$$S[n] = \Theta(V[n] - \vartheta_{\text{th}}), \quad (14)$$

$$V[n] = V[n](1 - S[n]) + V_{\text{reset}}, \quad (15)$$

where the the neuron will reset its membrane potential to $V_{\text{reset}} < \vartheta_{\text{th}}$ at time n once it trigger a spike $S[n]$. As described in Sec. 3.3, we choose to guide the membrane potential without reset stage (Eq. 15) $U[n]$ (or named no-reset membrane potential), instead of the normal membrane potential $V[n]$. The no-reset membrane potential is defined similarly as $V[n]$:

$$U[n] = \beta U[n - 1] + \gamma I[n], \quad (16)$$

$$S[n] = \Theta(U[n] - \vartheta_{\text{th}}), \quad (17)$$

but the neuron does not reset its membrane potential in this condition. The reason behind this choice is that the reset process will affect the guidance of the $V[n]$. Specifically, suppose we expect the neuron to fire a spike at $n + 1$, but if a spike just occurs at time n , the membrane potential $V[n]$ will reset to V_{reset} , consequently leading to a small value for $V[n + 1]$. Based on the SPA-Loss function, $V[n + 1]$ would then be guided by a large expected membrane potential value (above the threshold). However, this would incorrectly guide the membrane potential after resetting to the desired membrane potential, rather than guiding the true membrane potential as intended. Therefore, we choose to use the no-reset membrane potential $U[n]$ to effectively guide the spiking neuron to fire spike at the specified location.

H Proof of Proposition

Proposition 1. *Suppose the input event cell sequence has length N , desired spiking time is n^* ($n^* \in \{0, 1, \dots, N\}$), the membrane potential at time n^* satisfying the constraints:*

$$V_{th} \leq U[n^*] \leq \max(\beta V_{th} + \gamma I[n^*], V_{th}), \quad (18)$$

where $I[n^*]$ is the input synaptic current from Eq.1. Then the spiking neuron fires a spike at time n^* and does not excite spikes at neighboring moments.

Proof. Here we consider two conditions that affect the spiking state at moment n^* :

- (1) Membrane potential at time n^* is too small to emit a spike.
- (2) Membrane potential at time n^* is too large, affecting neighboring moment spiking states.

To satisfy the condition (1), we only need to guide the membrane potential $U[n^*]$ to reach the threshold V_{th} at time n^* , thus the upper bound of $U[n^*] = V_{th}$. In condition (2), we need to consider the state of the membrane potential at $n^* - 1$ and $n^* + 1$. We first exhibit the accumulation rules of membrane potential:

$$U[n^*] = \beta U[n^* - 1] + \gamma I[n^*] \quad (19)$$

However, if $U[n^*]$ is too large, this may cause the membrane potential $U[n^* - 1]$ to exceed the threshold and occur spike generation prematurely, and then the membrane potential will immediately drop to a reset value (Eq. 15). This will leave the membrane potential $U[n^*]$ at a very low value, making it difficult to trigger a spike. Hence, we should control the membrane potential not to exceed the threshold value at moment $n^* - 1$:

$$U[n^* - 1] = \frac{(U[n^*] - \gamma I[n^*])}{\beta} \leq V_{th} \quad (20)$$

$$\Rightarrow U[n^*] \leq \beta V_{th} + \gamma I[n^*] \quad (21)$$

Thus, the upper bound of $U[n^*] = \beta V_{th} + \gamma I[n^*]$. However, if the leaky factor β is small, there exists a possibility that $\beta V_{th} + \gamma I[n^*] \leq V_{th}$, thus we set the upper bound of $U[n^*]$ as $\max(\beta V_{th} + \gamma I[n^*], V_{th})$.

Next, we consider whether the spike at time n^* affects the pulse state at time $n^* + 1$. Since the neuron at the n^* moment has already completed the spike generation before accumulating $U[n^* + 1]$. Therefore the membrane potential at n^* does not affect the neuronal state at $n^* + 1$. In sum, if the membrane potential satisfies: $V_{th} \leq U[n^*] \leq \max(\beta V_{th} + \gamma I[n^*], V_{th})$, the spiking neuron fires a spike at moment n^* and does not excite spikes at neighboring moments.

I More Explanations in Spiking Position-aware Loss

I.1 Details in Membrane Potential-driven Loss

In Sec. 3.3.1, we explore the range of the expected membrane potential $U[n^*]$ and ensure its rationality by proposition 1 (proved in Appendix. H). To understand the setting of the membrane potential more easily, we visualize the boundary cases in Fig. 8.

The lower bound case means that the membrane potential at the desired index should be at least V_{th} to activate a spike, and the upper bound case guides the membrane potential to exceed the threshold but prevents generating spikes in the previous time step. Hence, the desired membrane potential should be bounded in $[V_{th}, \max(\beta V_{th} + \gamma I[n^*], V_{th})]$ and $\alpha \in [0, 1]$ (Eq. 8) balances the desired membrane potential $U[n^*]$ between U_{lower} and U_{upper} .

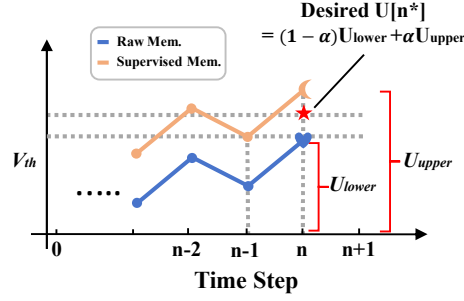


Figure 8: Visualization of the boundary cases when controlling the desired membrane potential, where the ‘heart-like’ point denotes the lower bound case and the ‘moon-like’ point denotes the upper bound case.

I.2 Details in Linear-assuming loss

As described in Sec. 3.3.2, suppose we expect the SNN to trigger a spike at a later time, if there exists a hill effect, the earlier membrane potential always reaches the excited state sooner and turns the neuron into the resting state to suppress the spiking generation at later moments. To address this challenge, we expect the later membrane potential to satisfy: (1) the later membrane potential should be larger than the current membrane potential to reverse the hill effect, and (2) the later membrane potential should exceed the threshold to fire a spike. Hence, we assume that the membrane potential in this condition should increase monotonically with the time step, as illustrated in Fig. 9.

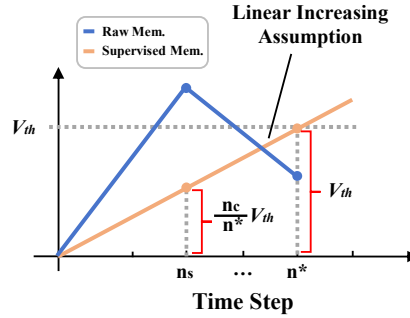


Figure 9: Visualization of our expected linearly increasing membrane potential.

Then we use the LA-Loss to supervise the membrane potential at n_c to reach $\frac{n_c}{n^*}V_{th}$ in order for the latter membrane potential at the n^* to reach V_{th} , satisfying both (1) and (2).

I.3 Details of Dynamic Hyperparameter Tuning

To deeply explore the control of the hyperparameter α , we first analyze the effect of α in Eq. 8. When α reaches its maximum value (*i.e.* $\alpha = 1$), the desired membrane potential evolves to U_{upper} , which corresponds to a situation where the membrane potential just reaches the threshold at the previous moment (in Fig. 8). That is, as α increases, the previous moment is more likely to generate a spike, driving the spiking time earlier. Recalling observation 2, taking a large- α scenario as an example, the SNN fails to spike at a later time step due to the alpha being too large, limiting the neuron’s ability to spike later. Hence, we expect the α to decrease to allow the desired index to decrease as well; similarly for small α scenarios. To summarize, we hope the α is updated in the same direction as the desired index, we update α by setting:

$$\alpha.\text{grad} = \left\| \sum_i^{N_s} (n^{*i} - n_c^i) / N_s \right\|_2^2, \quad (22)$$

$$\alpha \leftarrow \alpha - 2 \cdot \eta \sum_i^{N_s} (n^{*i} - n_c^i) / N_s, \quad (23)$$

where grad denotes the gradient of α and $'$ denotes derivative operation. The explanations of other math symbols can be found in the main text and Alg. 2.

J Visualization of the SNN Training Process

To validate whether the proposed feedback-update strategy can serve a guiding role during the initial stages of training, we have visualized the training process of the SNN and presented it in Figure 1 of the supplementary PDF rebuttal file. As anticipated, the training of the SNN exhibited fluctuations during the initial training stage, which might be attributed to the instability of the event quality obtained from dynamic slicing at this early phase. However, as training progressed, the loss of the SNN gradually stabilized and decreased, converging towards a desired outcome. Correspondingly, the slicing times progressively converged towards the desired spiking index. Therefore, although initial exploration may require several steps, our proposed training method is capable of offering effective guidance.

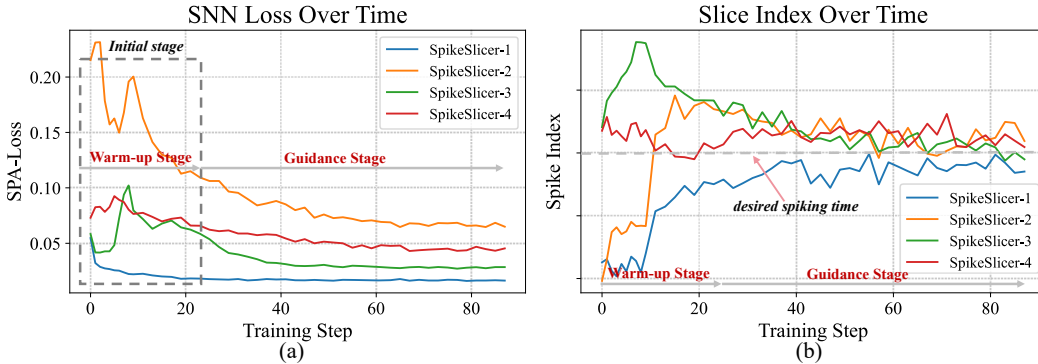


Figure 10: **Visualization of the SNN Training Process.** As shown in the gray box in (a), the SNN loss fluctuates considerably during the initial training stage. This is due to the instability in the quality of dynamic sliced events during the early phase (warm-up stage). However, as training progresses, the SNN loss gradually stabilizes and converges (guidance stage), with the corresponding slicing time (b) also converging towards the desired spike index.

K Implementation Details in Toy Experiments

To validate the effectiveness of SPA-Loss, we set up the toy task:

Input N randomized event cells, expect the SNN to slice at a specified time step n^ and there exists a certain probability of interfering with SNN to slice at other time steps.*

In detail, we set the time step within the range $[1, 30]$ and the max number of iterations as 800. We utilize a lightweight convolutional SNN with random initialization. We compare our proposed SPA-Loss function with common mean square error (MSE-Loss) and cross-entropy loss (CE-Loss). The results in Fig. 4 demonstrate that our proposed SPA-Loss can guide the SNN to spike at the desired timestep accurately with fewer convergence iterations than standard MSE-Loss and CE-Loss, paving the way for complex event-based vision tasks in Sec. 4.2. More experiments on the beginner’s arena are shown in Appendix N.

L Implementation Details in Event-based Task

We adopt a spiking neural network with structure: $\{16C3-GN-IF-AvgP2-32C3-GN-IF-AvgP2-64C3-GN-IF-AdaP2-LN-IF-LN-IF\}$, which consists of three convolutional layers and two linear layers, no residual block or attention are used. $\{i\}C\{j\}$ denotes a convolutional layer with the output channel i and the kernel size j ; GN denotes group normalization; $AvgP\{k\}$ and $AdaP\{k\}$ mean the average pooling and adaptive pooling with kernel size k ; LN denotes the linear layer. We choose the IF neuron as the activation function. We adopt the SGD optimizer and set the initial learning rate as $1e-4$, along with the cosine learning rate scheduler. SNN models are trained with batch size 32. Each experiment is conducted in an NVIDIA 4090 GPU.

L.1 Event-based Object Tracking

Datasets. The FE108 dataset [3] is an extensive event-based dataset for single object tracking, including 21 different object classes and several challenging scenes, *e.g.*, low-light (LL) and high dynamic range (HDR). The event streams are captured by a DAVIS346 event-based camera, which equips a 346x260 pixels dynamic vision sensor (DVS). We choose 54 sequences for training ANNs, 22 sequences for training SNNs and the rest 32 sequences for testing.

Evaluation Metrics. To show the quantitative performance of each tracker, we utilize three widely used metrics: success rate (Suc.), precision rate (Prec), normed precision rate (N-Prec), and overlap precision (OP). These metrics represent the percentage of three particular types of frames. Success rate is the frame of that overlap between the ground truth and the predicted bounding box is larger than a threshold; Precision rate focuses on the frame of the center distance between ground truth and predicted bounding box within a given threshold; OP_{thres} represents SR with *thres* as the threshold. We employ the area under curve (AUC) to represent the success rate. The precision score is associated with a 20-pixel threshold.

Label Settings. In this paper, since the original event dataset only provided labels at fixed frame rates, we employed a linear interpolation method to obtain corresponding labels for each more refined event cell. For example, suppose that in the original event dataset, a sub-event stream E with a period of T (i.e., $t \in [t_1, t_2]$) has labels $\{l_{t_1}, l_{t_2}\}$ corresponding to moments t_1, t_2 , respectively. If the number of event cells in this interval is N , then each event cell represents the event with the time range of $\{[t_1, t_1 + \frac{T}{N}], [t_1 + \frac{T}{N}, t_1 + 2\frac{T}{N}], \dots\}$, and the label for each interval can be derived through linear interpolation using $\{l_{t_1}, l_{t_2}\}$ and the number of event cells N . Thus, predictions at any slicing interval have corresponding labels for supervised learning. To ensure fairness, all tracking experiments in this paper utilize the aforementioned method to process the event dataset.

L.2 Event-based Recognition

DVS-Gesture. The DVS-Gesture [33] dataset contains 11 hand gestures from 29 subjects under 3 illumination conditions, recorded by a DVS128.

N-Caltech101. The N-Caltech101 dataset [34] incorporates 8,831 event-based images, with a 180×240 resolution and 101 classes, generated from the original Caltech101 dataset through an event-based sensor.

DVS-CIFAR10. The DVS-CIFAR10 dataset [35] is an event-stream dataset designed for object classification. It consists of 10,000 event streams, created by converting the frame-based images from the CIFAR-10 dataset using an event-based sensor with a resolution of 128×128 pixels. This dataset presents an intermediate level of difficulty, featuring 10 distinct classes.

SL-Animals. The SL-Animal database [36] features DVS recordings of individuals performing sign language gestures representing various animals, captured as a continuous spike flow with very low latency. This dataset includes approximately 1100 samples from 58 subjects, each performing 19 different sign language gestures in isolation across various scenarios, offering a challenging evaluation platform for this emerging technology.

M Theoretical Energy Consumption Calculation

To calculate the theoretical energy consumption, we begin by determining the synaptic operations (SOPs). The SOPs for each block in the SNN can be calculated using the following equation:

$$\text{SOPs}(l) = fr \times T \times \text{FLOPs}(l) \quad (24)$$

where l denotes the block number in the SNN, fr is the firing rate of the input spike train of the block and T is the time step of the spike neuron. $\text{FLOPs}(l)$ refers to floating point operations of l block, which is the number of multiply-and-accumulate (MAC) operations. And SOPs are the number of spike-based accumulate (AC) operations.

To estimate the theoretical energy consumption of SNN, we assume that the MAC and AC operations are implemented on a $45nm$ hardware, with energy costs of $E_{MAC} = 4.6pJ$ and $E_{AC} = 0.9pJ$, respectively. According to [43, 44], the calculation for the theoretical energy consumption of SNN is

given by:

$$E_{\text{Diffusion}} = E_{MAC} \times \text{FLOP}_{\text{SNN}_{\text{Conv}}}^1 + E_{AC} \times \left(\sum_{n=2}^N \text{SOP}_{\text{SNN}_{\text{Conv}}}^n + \sum_{m=1}^M \text{SOP}_{\text{SNN}_{\text{FC}}}^m \right) \quad (25)$$

where N and M represent the total number of layers of Conv and FC, E_{MAC} and E_{AC} represent the energy cost of MAC and AC operation, $\text{FLOP}_{\text{SNN}_{\text{Conv}}}^1$ denotes the FLOPs of the first Conv layer, $\text{SOP}_{\text{SNN}_{\text{Conv}}}^n$ and $\text{SOP}_{\text{SNN}_{\text{FC}}}^m$ are the SOPs of n^{th} Conv and m^{th} FC layer, respectively.

N More Experiments on Beginner’s Arena

Problem Setup. To verify the accuracy of our proposed slicing method, we expect SNN can slice events at the specified time step. We set up two scenarios and only show the difficult task (II) in the main text:

- *Task (I): Input T identical event cells, expect the SNN to slice at a specified time step T^* .*
- *Task (II): Input T randomized event cells, expect the SNN to slice at a specified time step T^* , but there exists a certain probability of interfering with SNN to slice at other locations.*

Task (I) aims to verify whether our proposed slicing strategy can accurately locate the optimal point; To test the robustness of our method, task (II) simulates complex event stream processing with random inputs and adds random noise to affect the SNN with wrong labels.

Table 9: Results on simple event slicing tasks with SPA-Loss.

	Input Size	Time Steps	Parameter	Iterations to Convergence ↓
<i>Task (I)</i>	32×32	30	0.52M	75
	64×64	30	2.02M	81
<i>Task (II)</i>	32×32	100	0.52M	29
	64×64	100	2.02M	88

We adopt a lightweight SNN (0.25M/2.02M) for our experiments. T^* is randomly selected within range $[0, T]$. The experimental results presented are the average of the results obtained by setting up three random seeds. As shown in Tab. 9, SNN requires only a small number of iterations to converge to the specified slicing time based on the SPA-Loss. For the complex task with random inputs and disturbances in task (II), SNN can still converge fast and even faster to find the specified cut point compared with task(I). *This simple experiment demonstrates that SPA-Loss can effectively supervise the SNN pulsing at the specified location, which paves the way for experiments on adaptive event slicing in real scenarios.*

O Statistics of Dynamic Slicing (SpikeSlicer) vs. Fixed Slicing

In this section, we compare the statistics results of the resulting events sliced by different slicing methods.

Symbol Description: the total event stream E ; the resulting sliced sub-event stream list by SNN: $E_{\text{beef}} = [E_1^b, \dots, E_{N_1}^b]$; the resulting sliced sub-event stream list by fixed slicing method: $E_{\text{fixtime}} = [E_1^f, \dots, E_{M_1}^f]$.

In the tracking task, the average duration of each sub-event stream $E_k^b (k \in [1, N_1])$ is 65ms (corresponding to 13 event cells, and the duration of the event stream contained in each event cell is 5ms). The maximum duration of each sub-event stream is 100ms, and the minimum duration is 30ms, while for our comparison of the slicing-by-fixed-time approach, the duration of each sub-event stream $E_j^f (j \in [1, M_1])$ is fixed at 75ms. The following Tab. O and Fig. 11 show the specific statistics of adaptive slicing vs. fixed slicing:

P Statistics of Resulting Slicing Duration

To further verify the stability and effectiveness of the dynamic slicing method, we explore the results of our method by changing the number of event cells in the event recognition task. N_{cell} indicates

Method	Avg Cell Num	Var Cell Num	Avg Duration	Min Duration	25th Duration	75th Duration	Max Duration
SpikeSlicer	12.99	3.96	~ 65ms	25ms	50ms	80ms	100ms
Slice by fixed duration	15	0	75ms	//	//	//	//

Table 10: Statistic results of dynamic slicing method (our SpikeSlicer) and fixed slicing method.

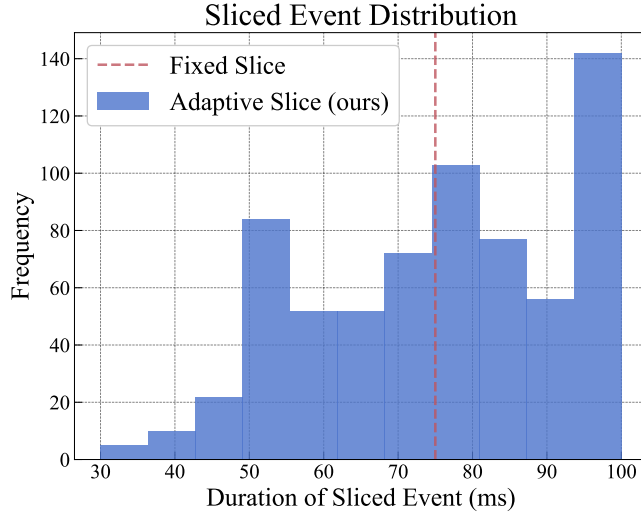


Figure 11: Visualization of sliced event distribution. Our method can return the sliced events with different durations, while the fixed method can only generate sliced events with fixed durations.

that an event stream is divided into N_{cell} event cells, and the larger the N_{cell} implies that the event stream is divided into more fine-grained event cell sequences that are capable of better represent the raw event stream (as mentioned in Sec. 3.1).

Calculation Process: Suppose the whole event stream (duration = T) is divided into 15 event cells, if the SNN is trained to sliced the event with 2.42 (average) event cells, which means that the sliced sub-event stream E_k^b contains event data which lasts a duration of $\frac{1}{15} * 2.42 * T = 16.13\%T$.

Table 11: Experiments of SpikeSlicer with different event cell numbers N . The resulting sliced event group always has a similar time interval in various N conditions.

N_{cell}	15	20	25
Avg Cell Num	2.42	3.15	4.77
Percentage of Duration	16.13%	15.75%	19.08%

The experimental results show that the percentage of the duration of each sub-event stream to the total event stream duration after the adaptive slicing is relatively stable, *i.e.*, the fineness of the event cell does not affect the event information contained in each sub-event stream after the slicing process, which proves the robustness and effectiveness of the dynamic slicing process of our method. We also illustrate the event percentage change during training in Fig. 12.

Q More Experiments with Latest Models

To enhance the credibility and robustness of our results, we have incorporated state-of-the-art models: Swin Transformer (SwinT [45]) and Vision Transformer (ViT [46]), to further validate the efficacy of our algorithm in event-based recognition tasks (Tab. 12):

Experiment Settings: We choose SwinT-small and ViT-small for comparisons on the DVSGesture dataset. Other settings are consistent with the main experiments.




Figure 12: Visualization of the percentage of event in various N_{cell} conditions during training. SpikeSlicer can always split and obtain sub-event groups with similar time intervals.

Table 12: Experiments of utilizing SpikeSlicer on latest recognition backbones.

Method	Random Slice	Fixed Slice	Ours
SwinT [45]	88.19	89.93	91.67(+1.74%)
ViT [46]	87.50	85.07	88.54(+3.47%)

Results demonstrate that the SpikeSlicer also yields performance improvement in recognition tasks with the latest backbones.

R More Experiments with Complex Neuromorphic Dataset

To further validate the proposed approach, we conduct experiments on more complex neuromorphic N-ImageNet [47]:

Table 13: Experiments of utilizing SpikeSlicer on complex dataset N-ImageNet [47]

Slice Method	Random Slice	Fixed Slice	Ours
ResNet-18	40.98	39.43	45.48(+6.05%)

S Impact Statement

This paper proposes an effective event processing method and also provides a novel SNN-ANN cooperation paradigm, aiming to inspire further research and development in energy-efficient and high-performance computing. We do not anticipate a direct negative impact from our work.

T Summary

To sum up, SpikeSlicer is designed as a **plug-and-play** algorithm for dynamic event stream slicing. It is benchmarked against baselines that employ fixed event stream slicing methods, proving the effectiveness of our method. In addition, our approach is versatile and can be applied in any event-based vision task, not limited to recognition or single object tracking scenarios.

Notably, SpikeSlicer also provides a brand-new **SNN-ANN cooperation paradigm**, where the SNN acts as an efficient, low-energy data processor to assist the ANN in improving downstream performance, injecting new perspectives and potential avenues of exploration.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading “NeurIPS paper checklist”,**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction of this paper clearly reflect the contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Please see Discussion and Conclusion.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: See Proposition 1 and its proof in Appendix H.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We have provided the experiment details in Appendix L.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The experiment details are provided in the paper. We will release the code soon.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.

- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The experiment details are provided in the Appendix L, where the network architecture, optimizer, batch number and other information are introduced.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We fixed the random seed.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Each experiment is conducted with an NVIDIA 4090 GPU. More details are provided in Appendix L.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: The paper is with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We provide the impact statement in Appendix S.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cited the used assets correctly.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not introduce any new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.