

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 EVENT2VEC: PROCESSING NEUROMORPHIC EVENTS DIRECTLY BY REPRESENTATIONS IN VECTOR SPACE

Anonymous authors

Paper under double-blind review

## ABSTRACT

Neuromorphic event cameras possess superior temporal resolution, power efficiency, and dynamic range compared to traditional cameras. However, their asynchronous and sparse data format poses a significant challenge for conventional deep learning methods. Existing solutions to this incompatibility often sacrifice temporal resolution, require extensive pre-processing, and do not fully leverage GPU acceleration. Inspired by word-to-vector models, we draw an analogy between words and events to introduce event2vec, a novel representation that allows neural networks to process events directly. This approach is fully compatible with the parallel processing and self-supervised learning capabilities of Transformer architectures. We demonstrate the effectiveness of event2vec on the DVS Gesture, ASL-DVS, and DVS-Lip benchmarks. A comprehensive ablation study further analyzes our method’s features and contrasts them with existing representations. The experimental results show that event2vec is remarkably parameter-efficient, has high throughput, and can achieve high accuracy even with an extremely low number of events. Beyond its performance, the most significant contribution of event2vec is a new paradigm that enables neural networks to process event streams as if they were natural language. This paradigm shift paves the way for the native integration of event cameras with large language models and multimodal models. Code, model, and training logs are provided in [https://anonymous.4open.science/r/event2vec\\_iclr-7B40](https://anonymous.4open.science/r/event2vec_iclr-7B40).

## 1 INTRODUCTION

Neuromorphic computing is an emerging research field that seeks to develop the next generation of artificial intelligence by emulating the brain’s principles (Mead, 1990). A significant advancement stemming from this paradigm is the event camera, a sensor inspired by the biological retina (Gallego et al., 2022). Prominent examples include the Dynamic Vision Sensor (DVS) (Lichtsteiner et al., 2008) and the Asynchronous Time-based Image Sensor (ATIS) (Posch et al., 2011). Unlike traditional cameras that capture synchronous frames, event cameras operate asynchronously, generating events in response to per-pixel brightness changes. This operational principle endows them with exceptionally high temporal resolution (on the order of microseconds), low power consumption, and a High Dynamic Range (HDR) exceeding 120 dB. This asynchronous operation results in a sparse stream of events, typically encoded in the Address-Event Representation (AER) format. An event is represented as a tuple  $(x, y, t, p)$ , composed of the pixel’s spatial coordinates  $(x, y)$ , a timestamp  $t$ , and a binary polarity  $p$  that indicates the direction of the brightness change.

Most contemporary deep learning models are designed to operate on dense, regularly structured, multi-dimensional tensors. This tensor-based paradigm is foundational to mainstream deep learning (LeCun et al., 2015) and is ubiquitously employed in modern scientific computing and machine learning frameworks, including NumPy (Harris et al., 2020a), TensorFlow (Abadi et al., 2016), and PyTorch (Paszke et al., 2019). Consequently, the sparse and asynchronous nature of event streams in the AER format is fundamentally incompatible with these tensor-based methods. To address this disparity, substantial research efforts have been devoted to converting events to dense representations, or designing new data and network structures to process the irregular events directly.

Existing methods primarily address the challenge of event encoding: how to effectively extract information from events and represent it for processing by neural networks. This challenge is analogous

054  
055  
056  
057  
058  
059  
060061  
062  
063  
Figure 1: Conceptual analogy between words and events. The illustration of the DVS 128 camera is  
adapted from Lichtsteiner et al. (2008).064  
065  
066  
067  
068  
069  
070  
to word encoding in natural language processing, a problem successfully addressed by foundational  
techniques such as word-to-vector (word2vec) (Mikolov et al., 2013). The word2vec model embeds  
each word into a fixed-length vector, enabling the relationships between words to be represented by  
mathematical operations between vectors. This vector representation approach is highly compatible  
with deep learning architectures and has become a foundational component of modern Natural Lan-  
guage Processing (NLP) models (Devlin et al., 2019a; Brown et al., 2020). We identify numerous  
parallels between words and events, as illustrated in Figure 1. The key similarities are as follows:071  
072  
073  
074  
075  
076  
077  
(1) **Each element is a composite of an index and a position.** In NLP, each word is assigned a  
unique index from a vocabulary, a conversion handled by a tokenizer; the indices in Figure 1,  
for instance, are generated by the Llama-3 tokenizer (Grattafiori et al., 2024). A word’s position  
is its sequential location within the sentence (e.g., the word “how” is at position 0 in “how  
are you”). Similarly, an event’s index is its spatial address, represented by the tuple  $(x, y, p)$ .  
Crucially, its position is not the sequence number, but its timestamp  $t$ , which marks its precise  
temporal location in the event stream.  
078  
079  
080  
081  
(2) **The set of possible indices is finite.** The vocabulary of a language, which forms the dictionary  
used in NLP, is finite. Likewise, an event camera has a limited set of possible event indices,  
defined by its sensor’s properties. For example, a DVS128 camera has  $2 \times 128 \times 128$  unique  
indices, corresponding to 2 polarities across a  $128 \times 128$  spatial resolution.  
082  
083  
084  
085  
(3) **The sequence exhibits a natural ordering.** Words in a sentence are arranged in a specific  
sequence that dictates meaning. Analogously, events are naturally ordered by their timestamps,  
reflecting the chronological progression of captured changes. This inherent temporal order is a  
key characteristic that distinguishes event data from unordered data structures like point clouds.  
086  
087  
088  
089  
090  
091  
092  
(4) **The meaning of an element is determined by its context.** A word can be polysemous; for  
instance, “transformer” can refer to a neural network architecture or a character in an animated  
series; its specific meaning is disambiguated by the surrounding text. An individual event merely  
signals a brightness change at a specific pixel and time, conveying little information in isolation.  
However, when viewed within a spatiotemporal stream, a sequence of events can delineate an  
object’s contour, thus giving a single event a higher-level meaning, such as being part of an  
edge. Therefore, the significance of an event is also fundamentally context-dependent.093  
094  
Inspired by word2vec, we propose event-to-vector (event2vec), an efficient spatio-temporal repre-  
095  
096  
097  
098  
099  
100  
101  
102  
103  
104  
105  
106  
107  
sentation for asynchronous events. Our contributions are as follows:

- By embedding events into a vector space, our method natively handles the sparse nature of the input stream, avoiding dense intermediate representations like event frames. This allows for efficient, GPU-accelerated processing with modern network architectures.
- We propose a parametric spatial embedding and a convolution-based temporal embedding method that captures neighborhood similarity—a task that is critical for accuracy but difficult for a standard embedding layer to learn.
- We validated our method on three classification benchmarks: DVS Gesture, ASL-DVS, and DVS-Lip. It achieved accuracy competitive with traditional methods while demonstrating remarkable parameter efficiency, throughput, and robustness, especially with a low number of events.

Beyond the performance metrics, the most significant contribution of event2vec is its ability to enable neural networks to process event streams in a manner analogous to natural language. Therefore,

108 state-of-the-art NLP architectures and methods—such as Transformer variants, model acceleration  
 109 algorithms, and generative self-supervised training—can be directly leveraged for event-based vi-  
 110 sion. Event2vec paves the way for the native integration of event cameras with large language  
 111 models and multimodal models.  
 112

## 113 2 RELATED WORK

### 115 2.1 DENSE REPRESENTATIONS OF EVENTS

117 Dense representations, derived from raw event streams, are fully compatible with conventional deep  
 118 learning methods. This is typically achieved by integrating events along the time axis to form dense  
 119 3D or 4D tensors, such as event frames (Liu & Delbruck, 2018), multi-channel images (Barchid  
 120 et al., 2022), voxel grids (Bardow et al., 2016), volumetric cubes (Cordone et al., 2022) and patches  
 121 (Sabater et al., 2023; Peng et al., 2023).

122 Specifically, event-to-frame methods accumulate events within discrete time intervals. These result-  
 123 ing frames can then be processed directly by standard neural networks. However, a significant draw-  
 124 back of these methods is the degradation or complete loss of the high temporal resolution inherent  
 125 to event data. This occurs because individual event timestamps are aggregated or quantized during  
 126 the conversion process. Furthermore, transforming the data into a dense representation negates the  
 127 inherent spatial sparsity of events. For instance, the generated frames often contain a substantial  
 128 number of zero-valued pixels. These pixels, while carrying no information, still incur significant  
 129 memory and computational overhead. While many methods use timestamps implicitly to define the  
 130 integration interval, some approaches explicitly leverage them to generate temporal weights (Zhu  
 131 et al., 2019; Gehrig et al., 2019). Finally, the conversion process itself can be computationally in-  
 132 tensive, introducing considerable latency that is often prohibitive for real-time applications (Rebecq  
 133 et al., 2019; Gallego et al., 2022).

### 134 2.2 IRREGULAR REPRESENTATIONS OF EVENTS

136 Conversely, methods for processing irregular representations aim to preserve the inherent sparsity  
 137 and asynchronicity of event data. This category includes SNNs (Maass, 1997; Roy et al., 2019),  
 138 Sparse Convolutional Networks (Sparse CNNs) (Messikommer et al., 2020; Santambrogio et al.,  
 139 2024), Graph Neural Networks (GNNs) (Bi et al., 2019; Schaefer et al., 2022), and point-based  
 140 methods (Yang et al., 2019; Sekikawa et al., 2019; Lin et al., 2023; Ren et al., 2025).

141 When deployed on neuromorphic hardware (Merolla et al., 2014; Davies et al., 2018), Spiking Neu-  
 142 ral Networks (SNNs) can process events in a naturally asynchronous event-driven manner. However,  
 143 on standard hardware, GPU-based simulations of SNNs produce dense tensor outputs, as the hard-  
 144 ware necessitates synchronous processing with discrete time-steps. Consequently, training SNNs  
 145 on conventional GPUs typically occurs in a synchronous fashion, leading to an unavoidable perfor-  
 146 mance gap between synchronous training and asynchronous inference (Yao et al., 2024; Du et al.,  
 147 2025). Moreover, the reliance on backpropagation-through-time renders the training process slow  
 148 and memory-intensive. Sparse CNNs leverage the inherent sparsity of event data, achieving a theo-  
 149 retically low number of Floating-Point Operations (FLOPs). Nevertheless, the architecture of stand-  
 150 ard GPUs is not optimized for the dynamic computations and unstructured memory access patterns  
 151 required for efficient sparse acceleration. Consequently, similar to SNNs, Sparse CNNs fail to fully  
 152 exploit the massive parallel processing capabilities of GPUs.

153 Event-based GNNs construct graphs from incoming events, an approach that effectively preserves  
 154 the spatio-temporal relationships between them. Since empty regions with no event activity do not  
 155 generate graph nodes, the data’s sparsity is well-utilized. Their main disadvantage lies in the need  
 156 for careful hyper-parameter tuning, such as the event downsampling rate and neighborhood radius  
 157 for graph construction. Additionally, functioning as low-pass filters (Nt & Maehara, 2019), GNNs  
 158 are susceptible to the over-smoothing problem (Zhou et al., 2020), which limits their ability to form  
 159 deep architectures comparable to modern CNNs and Transformers (Vaswani et al., 2017). Point-  
 160 based methods treat events from event cameras as analogous to point clouds from Light Detection  
 161 and Ranging (LiDAR) sensors. A fundamental limitation of most point cloud models is their permu-  
 162 tation invariance, which necessitates treating the input as an unordered set. Consequently, the event  
 163 timestamp is typically relegated to being an additional positional coordinate, thereby discarding the

162 crucial causal ordering of events. To manage the data volume, these methods often employ classic  
 163 point cloud pre-processing techniques like farthest point sampling, which further increases latency.  
 164

### 165 3 METHODS

#### 167 3.1 REPRESENTING EVENTS IN A VECTOR SPACE

169 Leveraging the strong analogy between words and events, we propose a method for representing  
 170 events within a vector space, which we term event-to-vector (event2vec). An event, generated by a  
 171 camera with a spatial resolution of  $H \times W$ , is represented as a tuple  $(x, y, t, p)$ . For our embedding,  
 172 we treat the triplet  $(x, y, p)$  as the spatial coordinate and the timestamp  $t$  as the temporal coordinate.  
 173 The general formulation for the event2vec embedding is defined as:

$$174 \mathbf{v} = \mathbf{v}_s + \mathbf{v}_t = \text{Embed}_s(x, y, p) + \text{Embed}_t(t), \quad (1)$$

176 where  $\mathbf{v} \in \mathbb{R}^D$  is the resulting  $D$ -dimensional embedding vector,  $\mathbf{v}_s = \text{Embed}_s(x, y, p) \in \mathbb{R}^D$  is  
 177 the spatial embedding module, and  $\mathbf{v}_t = \text{Embed}_t(t) \in \mathbb{R}^D$  is the temporal embedding module. As  
 178 shown in Eq. 1, this method fuses spatial and temporal information through addition. This additive  
 179 fusion strategy is directly inspired by the positional encoding mechanism prevalent in Transformers.

#### 180 3.2 SPATIAL EMBEDDING

182 A straightforward approach for the spatial embedding module is to adapt the standard embedding  
 183 layer from NLP, which is efficiently implemented as a look-up table:

$$185 \mathbf{v}_s = \text{Embed}_s(x, y, p) = \mathbf{W}_s[p \cdot H \cdot W + y \cdot W + x], \quad (2)$$

186 where  $\mathbf{W}_s \in \mathbb{R}^{(2 \cdot H \cdot W) \times D}$  is the learnable embedding matrix and  $D$  is the embedding size. This  
 187 method maps each unique spatial coordinate to a distinct row index in the embedding matrix  $\mathbf{W}_s$ .

188 However, this standard embedding layer imposes no inductive bias on the relationship between  
 189 indices, compelling the model to learn all spatial relationships from data alone. In a tokenizer,  
 190 a word’s index is a non-semantic identifier, whose order is primarily determined by the token’s  
 191 frequency in the training corpus. Consequently, the words at indices  $i$  and  $i + 1$  share no inherent  
 192 semantic similarity. This assumption does not hold for event coordinates. Images are continuous  
 193 two-dimensional functions (Gonzalez, 2009). Spatially adjacent pixels are known to exhibit strong  
 194 correlation. Therefore, an effective spatial embedding should incorporate this locality bias, ensuring  
 195 that events with close coordinates yield similar embedding vectors:

$$196 \text{Embed}_s(x + \Delta x, y + \Delta y, p) - \text{Embed}_s(x, y, p) \approx \mathbf{0}, \quad (3)$$

197 for small coordinate perturbations  $[\Delta x, \Delta y]$ , e.g.,  $[\Delta x, \Delta y] = [1, 0]$ .

199 The standard embedding in Eq. 2 fails to account for this crucial spatial relationship, which can  
 200 impede the learning process. To solve this issue, we propose an elegant parametric algorithm for  
 201 generating the embedding matrix  $\mathbf{W}_\phi$  by a neural network  $\phi$ . To systematically enumerate all spatial  
 202 coordinates within a  $P \times H \times W$  volume (where  $P = 2$  represents the two polarities), we first  
 203 establish a linear index sequence  $\mathbf{c} = [0, 1, \dots, P \cdot H \cdot W - 1]$ . This sequence is then decomposed into  
 204 three probe tensors,  $\mathbf{x}_c$ ,  $\mathbf{y}_c$ , and  $\mathbf{p}_c$ , which correspond to the coordinates along the width, height, and  
 205 polarity dimensions, respectively. The transformation is defined as follows:  $\mathbf{x}_c = \mathbf{c} \pmod{W}$ ,  $\mathbf{y}_c =$   
 206  $\lfloor \frac{\mathbf{c}}{W} \rfloor \pmod{H}$ ,  $\mathbf{p}_c = \lfloor \frac{\mathbf{c}}{WH} \rfloor$ . Finally, these probe tensors are passed through  $\phi$ , which outputs  
 207 the complete embedding matrix  $\mathbf{W}_\phi = \phi(\mathbf{x}_c, \mathbf{y}_c, \mathbf{p}_c)$ . By substituting the parametrically generated  
 208 matrix  $\mathbf{W}_\phi$  into the look-up mechanism of Eq. 2, we establish a direct equivalence for any given  
 209 event coordinate  $(x, y, p)$ :

$$210 \mathbf{W}_\phi[p \cdot H \cdot W + y \cdot W + x] = \phi(x, y, p). \quad (4)$$

211 Crucially, the parametric network  $\phi$  is designed to be a continuous and differentiable function. This  
 212 property allows us to formally analyze the relationship between neighboring embeddings using a  
 213 first-order Taylor series expansion:

$$215 \phi(x + \Delta x, y + \Delta y, p) - \phi(x, y, p) \approx \frac{\partial \phi}{\partial x} \Delta x + \frac{\partial \phi}{\partial y} \Delta y + o(\|\Delta\|), \quad (5)$$

216 where  $o(\|\Delta\|)$  represents higher-order remainder terms. As Eq. 5 illustrates, for small perturbations  
 217  $[\Delta x, \Delta y]$ , the difference between the embeddings is proportional to the gradient of  $\phi$ . Consequently,  
 218 as the perturbations approach zero, this difference vector also approaches zero:

$$219 \quad 220 \quad \phi(x + \Delta x, y + \Delta y, p) - \phi(x, y, p) \approx \mathbf{0}. \quad (6)$$

221 In this manner, the use of a continuous parametric network  $\phi$  inherently embeds the desired neighbor-  
 222 hood semantics, or spatial inductive bias, directly into the embedding matrix. This approach  
 223 elegantly satisfies the condition outlined in Eq. 3.

### 225 3.3 TEMPORAL EMBEDDING

227 Timestamps, which denote the occurrence time of events, serve a function analogous to positional  
 228 indices in a sentence. In modern NLP models, relative positional encoding methods (Press et al.,  
 229 2021; Su et al., 2024) are increasingly favored over absolute methods, such as sinusoidal encoding  
 230 (Vaswani et al., 2017) or learnable absolute positional embeddings (Devlin et al., 2019b).

231 However, directly applying these relative positional encoding techniques to event timestamps is  
 232 ill-suited. Such methods are fundamentally designed for discrete and uniformly spaced indices,  
 233 whereas event timestamps are continuous and inherently non-uniform. To address this discrepancy,  
 234 we propose learning the temporal embedding directly from the differences between consecutive  
 235 timestamps using a neural network.

236 Specifically, the temporal embedding module is implemented as a stack of convolutional layers. For  
 237 the  $i$ -th event in an event stream, the input to this module is the first-order temporal difference,  
 238  $t[i+1] - t[i]$ . This design offers several distinct advantages:

- 239 (1) **Time-Shift Invariance:** By operating on relative temporal distances, the embedding becomes  
 240 inherently invariant to absolute shifts in time.
- 242 (2) **Generalization:** It functions as a variant of position-wise learnable encoding but circumvents  
 243 the length generalization problem by accepting continuous values as input, rather than being  
 244 tied to a fixed vocabulary of discrete positions.
- 245 (3) **Contextual Consistency:** The convolutional operations allow the temporal embedding for an  
 246 event to be influenced by the timing of its immediate neighbors, thereby reinforcing the principle  
 247 of neighborhood semantics in the time domain.

### 248 3.4 EVENT SAMPLING AND AGGREGATION

250 Raw event streams often contain an extremely large number of events, with sequence lengths exhibiting  
 251 substantial variance. Furthermore, deep learning frameworks typically process data in batches,  
 252 which requires that all tensors within a single batch have uniform dimensions. Consequently, it is  
 253 necessary to sample or aggregate events from each stream to a fixed-length sequence of size  $L$ .

254 In this paper, we primarily use two methods. The first is uniform random sampling. We find that  
 255 this straightforward method works well in most cases and is extremely computationally efficient.  
 256 However, a significant limitation of random sampling is the substantial information loss incurred  
 257 by discarding the majority of the events, leading to suboptimal accuracy in complex tasks. Our  
 258 second method addresses this by leveraging k-means clustering to aggregate the entire event stream  
 259 into  $L$  representative clusters. Specifically, the clustering process is performed independently on the  
 260 two event polarities to preserve their distinct information channels. After clustering, the centroid of  
 261 each cluster is treated as a representative event. Its timestamp is taken directly from the centroid's  
 262 temporal coordinate, while its spatial coordinates are quantized to the nearest integers. Furthermore,  
 263 we compute an intensity factor,  $\rho$ , equal to the number of raw events belonging to that cluster. This  
 264 intensity factor then modulates the corresponding spatial embedding vector, effectively weighting  
 265 the representation by its event density. This approach ensures that information from every event  
 266 contributes to the final representation.

267 In summary, the final event2vec representation for a sequence of  $L$  events is a tensor  $\mathbf{V} \in \mathbb{R}^{L \times D}$ .  
 268 The embedding for the  $i$ -th event in this sequence,  $\mathbf{V}[i]$ , is formulated as:

$$269 \quad \mathbf{V}[i] = \rho[i] \cdot \text{Embed}_s(\mathbf{x}[i], \mathbf{y}[i], \mathbf{p}[i]) + \text{Embed}_t(\Delta\mathbf{t})[i], \quad (7)$$

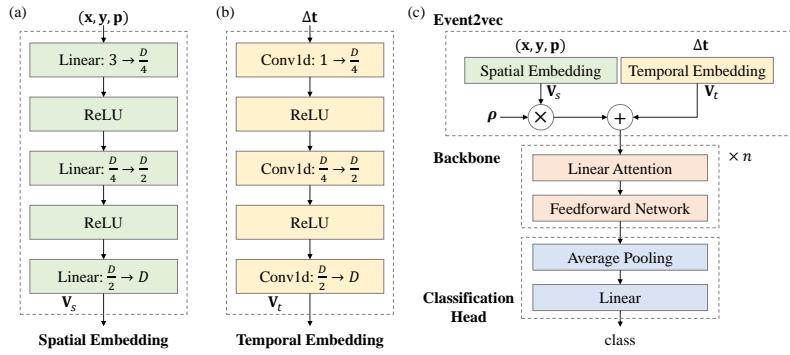


Figure 2: The network architecture for event classification using the event2vec representation.

where  $\rho$ ,  $\mathbf{x}$ ,  $\mathbf{y}$ ,  $\mathbf{p}$ , and  $\mathbf{t}$  are vectors representing the intensity factors, spatial coordinates, and timestamps for the sequence of  $L$  events. The vector of temporal differences is defined as  $\Delta\mathbf{t}[i] = \mathbf{t}[i+1] - \mathbf{t}[i]$ , with the final value set to zero. For a native event,  $\rho[i]$  is 1, while for a cluster event, it represents the number of raw events aggregated into that cluster.

### 3.5 NETWORK STRUCTURE

In this paper, we employ the Transformer architecture, leveraging its core strengths: the ability to efficiently process sequences in parallel and capture long-range dependencies. These characteristics make it particularly well-suited for the sequential representations generated by event2vec.

As shown in Figure 2(a), the spatial embedding module  $\phi$  consists of a stack of linear layers. It gradually increases the number of features from 3 to  $D$ . Layer Normalization (Ba et al., 2016) layers are also inserted to stabilize training but omitted from the figure for clarity. The temporal embedding module has a similar structure to the spatial embedding module, except that it uses depth-wise convolutional layers with a kernel size of 3 and a stride of 1, shown in Figure 2(b).

We employ the Forgetting Transformer (Lin et al., 2025) as the linear attention in Figure 2(c). It is important to recognize that linear attention is fundamentally equivalent to RNNs (Katharopoulos et al., 2020), operating with fixed-size hidden states. To enhance the learning capability for extremely long event sequences, we extend the linear attention model to a parameter-shared bi-directional formulation. Further details are provided in Appendix A.2. A linear attention module and a feedforward network, consisting of two linear layers, together constitute a single backbone block. As illustrated in Figure 2(c), the full backbone is composed of  $n$  such blocks stacked sequentially. For the classification head, we employ an average pooling layer to aggregate features across all positions in the sequence.

## 4 EXPERIMENTS

To validate the event2vec representation, we conduct a series of experiments on classification tasks using three neuromorphic datasets: DVS Gesture (Amir et al., 2017), ASL-DVS (Bi et al., 2019), and DVS-Lip (Tan et al., 2022). In this section, results are reported in the format  $a \pm b$ , representing the mean and standard deviation, respectively. For experiments that involve random sampling, these statistics are computed over 10 independent runs on the test set. Experimental details are provided in Appendices A.3 to A.8.

### 4.1 COMPARISON BETWEEN REPRESENTATIONS

**Accuracy and Parameter Efficiency** Table 1 compares the accuracy and model parameters of event2vec with those of other representations across the three datasets. Our models for DVS Gesture and ASL-DVS are trained directly on randomly sampled events. For DVS-Lip, our model first undergoes self-supervised pre-training (refer to Appendix A.5) on cluster events. We then report the fine-tuning accuracy on both randomly sampled events and cluster events. Our method achieves accuracy comparable to that of other leading representations while demonstrating exceptional parameter efficiency.

324  
325  
Table 1: Model performance and size comparison on neuromorphic datasets

326 Dataset	327 Method + Representation	328 Accuracy (%)	329 Params (MB)
327 DVS Gesture	Sparse GRU + Frame (Subramoney et al., 2023)	97.80	4.8
	SNN + Frame (Yao et al., 2023)	<b>98.23</b>	6.5
	328 FARSE-CNN + Window Slicing (Santambrogio et al., 2024)	96.6	10.79
	329 Event MAE + Point Cloud (Sun et al., 2025)	97.75	Unknown
330 ASL-DVS	331 Linear Attention + Event2vec (4096 Randomly Sampled Events)	97.57 $\pm$ 1.31	<b>0.52</b>
	331 GNN,CNN + Graph (Bi et al., 2019)	90.10	19.46
	332 GNN,Transformer + Image,Voxel Graph (Yuan et al., 2023)	99.60	220.3
333 DVS-Lip	332 Linear Attention + Event2vec (512 Randomly Sampled Events)	<b>99.91<math>\pm</math>0.05</b>	<b>0.27</b>
	333 ResNet-18,BiGRU + Frame (Tan et al., 2022)	60.3	72.1
	334 Spiking ResNet18,BiGRU + Frame (Dampfhofer & Mesquida, 2024)	<b>75.3</b>	58.6
	335 Linear Attention + Event2vec (1024 Randomly Sampled Events)	70.62 $\pm$ 1.55	<b>18.3</b>
	335 Linear Attention + Event2vec (1024 Cluster Events)	75.14	<b>18.3</b>

336  
337  
Table 2: Comparative analysis of pre-processing latency for event representations on DVS Gesture

338 Representation	339 Hyper-Parameter	340 Total Pre-processing Time (s)
339 Random Sampling (used in event2vec)	4096 events	<b>1.24<math>\pm</math>0.04</b>
340 Frame (Yao et al., 2023)	16 frames	4.44 $\pm$ 0.03
341 Graph (Schaefer et al., 2022)	Radius=5, 32 neighbors, 10000 samples	6.37 $\pm$ 0.06
342 Voxel Grid (Zhu et al., 2019)	16 bins	7.70 $\pm$ 0.15
343 Point Cloud (Sun et al., 2025)	64 patches $\times$ 32 points	47.56 $\pm$ 1.50
344 K-Means (used in event2vec)	1024 clusters	119.51 $\pm$ 0.55
344 Window Slicing (Santambrogio et al., 2024)		230.79 $\pm$ 2.46

345  
346  
Table 3: Benchmark of throughput and single event stream inference latency on DVS Gesture

347 Method	348 Throughput (samples/s)	349 Single Stream
	348 Training	349 Inference
349 FARSE-CNN + Window Slicing (Santambrogio et al., 2024)	6.65 $\pm$ 0.85	23.94 $\pm$ 9.21
350 AEGNN + Graph (Schaefer et al., 2022)	111.06 $\pm$ 6.68	1271.75 $\pm$ 2.69
351 SNN + Frame (Yao et al., 2023)	202.88 $\pm$ 3.69	234.00 $\pm$ 7.61
352 Sparse GRU + Frame (Subramoney et al., 2023)	417.00 $\pm$ 0.72	472.82 $\pm$ 18.40
352 <b>Event2vec + Randomly Sampling</b>		<b>1030.02<math>\pm</math>56.36</b> <b>2884.56<math>\pm</math>283.89</b>
		11.44 $\pm$ 0.69

353  
**Event Pre-processing Time** Table 2 benchmarks the total pre-processing time required by different  
354 representations to process all 1,176 event streams of the DVS Gesture training set. Hyper-parameters  
355 were set to the values specified in their respective original publications, where available. For others,  
356 we adopted commonly used values; for instance, graph construction was constrained to an average  
357 node degree between 10 and 20 to ensure balanced connectivity, and the number of frames and bins  
358 of the voxel grid was set to 16, a common configuration (Zhou et al., 2024). The results indicate that  
359 event2vec with random sampling exhibits minimal pre-processing time. While employing K-Means  
360 clustering increases pre-processing time, it still remains faster than window slicing.

361  
**Throughput and Latency** Table 3 compares the throughput and single event stream inference la-  
362 tency of FARSE-CNN, AEGNN, SNN, Sparse GRU and our event2vec model on the DVS Gesture  
363 dataset. Our event2vec model exhibits extremely high training and inference throughput, primarily  
364 due to its full compatibility with a highly optimized linear attention library that can fully leverage  
365 GPUs for acceleration. The FARSE-CNN model exhibits the lowest throughput. This is attributable  
366 to its reliance on Sparse CNN, which receives limited acceleration from unstructured sparsity on  
367 GPUs, and the inclusion of a recurrent structure that further constrains its speed.

368  
4.2 ABLATION EXPERIMENTS

370  
**Embedding Comparison** We conducted an ab-  
371 lation study on the DVS Gesture dataset to  
372 evaluate the accuracy contributions of differ-  
373 ent components, as detailed in Table 4. We  
374 tested various combinations of spatial embed-  
375 ding methods (standard (Eq. 2) vs. paramet-  
376 ric (Eq. 4)) and temporal embedding modules  
377 (sinusoidal embedding on  $t$  vs. convolutional  
embedding on  $\Delta t$ ). The combination of the standard embedding with our convolutional temporal

378  
Table 4: Ablation analysis of embeddings on the  
379 DVS Gesture dataset

380 Embedding	381 Spatial	382 Temporal	383 Accuracy (%)
381 Standard	382 Conv( $\Delta t$ )	383 91.18 $\pm$ 3.70	
382 Standard	383 Sin( $t$ )	384 93.16 $\pm$ 2.19	
383 Parametric	384 Sin( $t$ )	385 96.56 $\pm$ 1.46	
384 Parametric	385 Conv( $\Delta t$ )	386 97.57 $\pm$ 1.31	

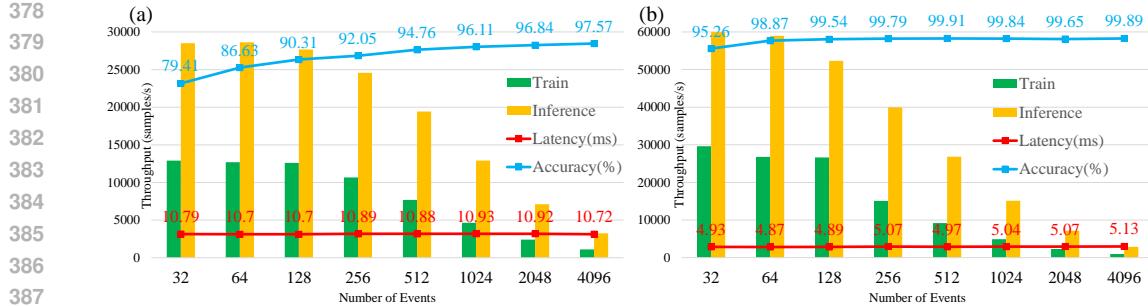


Figure 3: Effect of number of events on models for (a) DVS Gesture (b) ASL-DVS.

embedding (Standard +  $\text{Conv}(\Delta t)$ ) yields the lowest accuracy. We attribute this to the standard embedding layer’s lack of inductive bias, which prevents it from effectively learning neighborhood semantics and subsequently hinders the convolutional temporal encoder. Consequently, when using our parametric embedding, the convolutional encoder achieves the highest accuracy. It is worth noting that our parametric embedding consistently outperforms the standard version when paired with any temporal embedding, validating the effectiveness of incorporating neighborhood semantics.

**Effect of Event Numbers** Processing fewer events results in lower resource consumption, which is always desirable in event-based applications. We benchmark the impact of varying the number of randomly sampled events ( $L$ ) on several key metrics: training/inference throughput, single event stream inference latency, and accuracy on the DVS Gesture and ASL-DVS datasets. As illustrated in Figure 3, the performance trends are consistent across both datasets. For small values of  $L$ , throughput decreases only marginally as  $L$  increases. This is attributable to the CUDA kernel launch overhead, which dominates the computation time, rendering the actual kernel execution time negligible in this regime. As  $L$  grows larger, the kernel execution time becomes the primary bottleneck, causing a more pronounced decrease in throughput. Notably, while  $L$  increases, the throughput decreases approximately inversely proportionally, a finding consistent with the  $\mathcal{O}(L)$  complexity of linear attention. The single event stream inference latency for both models increases only slightly with  $L$ , further indicating that kernel launch overhead, rather than execution time, remains the dominant factor. In terms of accuracy, the general trend shows improvement as  $L$  increases. Surprisingly, both models maintain a reasonable level of accuracy even with a very small number of events (e.g.,  $L = 32$ ), demonstrating their robustness to sparse inputs.

We also compare our method with the sophisticated sampling techniques from Araghi et al. (2025), which use a voxel grid representation (Figure 4). The results highlight the inherent effectiveness of event2vec: when paired with simple random sampling, it consistently outperforms the voxel grid representation, even when the latter employs more complex, meticulously designed sampling strategies.

### 4.3 VISUALIZATION

**Neighborhood Semantics** To visually inspect the neighborhood semantics, we extract the spatial embedding weights from models trained on the DVS Gesture dataset with the parametric ( $\mathbf{W}_\phi$ ) and standard ( $\mathbf{W}_s$ ) embedding layers. For each coordinate  $(x, y, p)$ , its  $D$ -dimensional embedding vector is projected onto a 3-dimensional space using Principal Component Analysis (PCA). These 3D vectors are then interpreted as RGB color values and plotted at their corresponding  $(x, y)$  locations to form an image. Figure 5(a) visualizes the resulting images for polarity 0 (images for polarity 1 are provided in Appendix A.9). The image derived from  $\mathbf{W}_\phi$  displays smooth, continuous color gradients, akin to a color palette, indicating that spatially adjacent coordinates have semantically similar embeddings. In stark contrast, the image from  $\mathbf{W}_s$  resembles random noise, signifying a lack of learned spatial correlation.

**Polarity Similarity** An object’s edge moving across a pixel often triggers events of both polarities in close succession. We therefore hypothesize that the embeddings for opposite polarities at the same

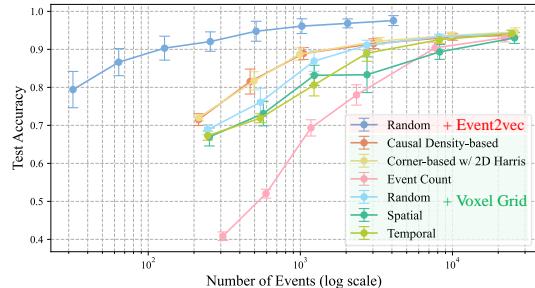


Figure 4: DVS Gesture: Accuracy vs. number of events against Araghi et al. (2025).

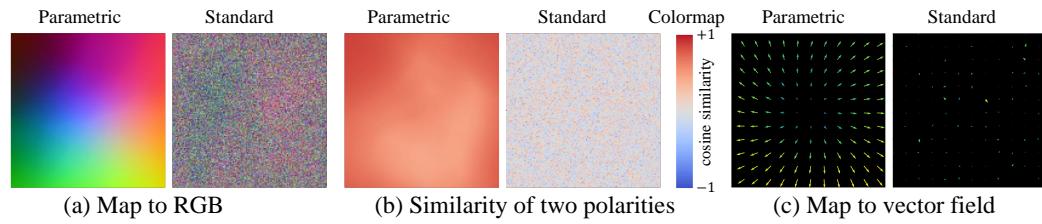


Figure 5: Visual comparison of the learned spatial embeddings: parametric vs. standard.

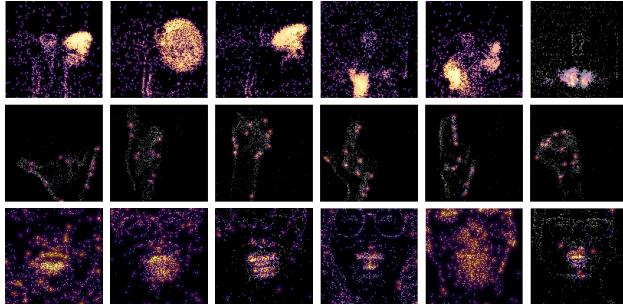


Figure 6: Event-level attention maps on samples from DVS Gesture (Row 1), ASL-DVS (Row 2), and DVS-Lip (Row 3).

spatial location should also be semantically related. To test this, we compute the cosine similarity between the embedding vectors of the two polarities at each coordinate. As shown in Figure 5(b), the parametric embedding captures this relationship, exhibiting distinct regions of high similarity. Conversely, the similarity map for the standard embedding is predominantly close to zero, indicating that it fails to learn this inter-polarity correlation.

**Vector Field Representation** We visualize the learned spatial manifold as a vector field. The  $D$ -dimensional embedding vectors are projected onto their first two principal components using PCA. These resulting 2D vectors are then visualized using a quiver plot, where each arrow represents the direction and magnitude of the vector at its spatial coordinate. Figure 5(c) illustrates the results. The vector field for the parametric embedding exhibits a coherent, laminar-like flow, revealing a smoothly structured semantic space. In contrast, the field for the standard embedding appears chaotic and turbulent, further confirming its inability to capture meaningful spatial relationships.

**Event-wise Attention** As event2vec is an event-wise representation, its attention mechanism can be visualized at a fine-grained, event-level resolution. Figure 6 displays attention heatmaps overlaid on the original event streams for DVS Gesture (row 1), ASL-DVS (row 2), and DVS-Lip (row 3). The visualizations reveal that the model correctly focuses on the hands in DVS Gesture, the finger joints and contours in ASL-DVS, and the lip region in DVS-Lip. However, consistent with the lower classification accuracy compared to the other two datasets, we also observe instances where the model incorrectly allocates significant attention to other facial features, such as the eyes and ears.

## 5 CONCLUSIONS

Neuromorphic event cameras introduce a paradigm shift in computer vision, presenting both unique opportunities and significant challenges. A central challenge has been reconciling their asynchronous, sparse nature with the synchronous, dense tensor-based architectures of deep learning. In this paper, we introduced event2vec, a novel representation that directly addresses this challenge by enabling neural networks to natively process asynchronous events. Our experimental results demonstrate that event2vec achieves accuracy competitive with established methods while offering compelling advantages in parameter efficiency, pre-processing overhead, throughput, and robustness across varying numbers of events. Beyond these performance metrics, the most significant contribution of event2vec is its conceptual alignment of event streams with the paradigm of natural language processing. This opens new avenues for research and application. By treating events as a sequential language, we can begin to explore novel applications by leveraging the sophisticated architectures developed for large language models.

486 REPRODUCIBILITY STATEMENT  
487

488 The experimental code, training logs, terminal output, and trained models for this paper are  
489 all provided together in the code repository [https://anonymous.4open.science/r/  
490 event2vec\\_iclr-7B40](https://anonymous.4open.science/r/event2vec_iclr-7B40). We have included detailed instructions in the repository, allowing  
491 readers to easily prepare the dataset and reproduce the experiments based on these instructions. All  
492 experiments in this paper fix the random number seeds of PyTorch, NumPy, and Python to 0 using  
493 the `seed_everything` function from PyTorch Lightning, in order to maintain reproducibility as  
494 much as possible.

495 REFERENCES  
496

497 Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu  
498 Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg,  
499 Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan,  
500 Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: a system for large-  
501 scale machine learning. In *Proceedings of the 12th USENIX Conference on Operating Systems  
502 Design and Implementation*, OSDI’16, pp. 265–283, USA, 2016. USENIX Association. ISBN  
503 9781931971331.

504 Arnon Amir, Brian Taba, David Berg, Timothy Melano, Jeffrey McKinstry, Carmelo Di Nolfo,  
505 Tapan Nayak, Alexander Andreopoulos, Guillaume Garreau, Marcela Mendoza, Jeff Kusnitz,  
506 Michael Debole, Steve Esser, Tobi Delbruck, Myron Flickner, and Dharmendra Modha. A low  
507 power, fully event-based gesture recognition system. In *Proceedings of the IEEE/CVF Conference  
508 on Computer Vision and Pattern Recognition*, pp. 7388–7397, 2017. doi: 10.1109/CVPR.2017.  
509 781.

510 Hesam Araghi, Jan van Gemert, and Nergis Tomen. Making Every Event Count: Balancing Data  
511 Efficiency and Accuracy in Event Camera Subsampling. In *IEEE/CVF Conference on Computer  
512 Vision and Pattern Recognition Workshops*, pp. 5044–5054, Los Alamitos, CA, USA, June 2025.  
513 IEEE Computer Society. doi: 10.1109/CVPRW67362.2025.00499.

515 Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint  
516 arXiv:1607.06450*, 2016.

517 Sami Barchid, José Mennesson, and Chaabane Djéraba. Bina-rep event frames: A simple and  
518 effective representation for event-based cameras. In *IEEE International Conference on Image  
519 Processing*, pp. 3998–4002. IEEE, 2022.

521 Patrick Bardow, Andrew J. Davison, and Stefan Leutenegger. Simultaneous optical flow and in-  
522 tensity estimation from an event camera. In *Proceedings of the IEEE Conference on Computer  
523 Vision and Pattern Recognition*, June 2016.

524 Yin Bi, Aaron Chadha, Alhabib Abbas, Eirina Bourtsoulatze, and Yiannis Andreopoulos. Graph-  
525 based object classification for neuromorphic vision sensing. In *IEEE/CVF International Confer-  
526 ence on Computer Vision*, pp. 491–501, 2019. doi: 10.1109/ICCV.2019.00058.

527 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
528 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel  
529 Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler,  
530 Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray,  
531 Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever,  
532 and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information  
533 Processing Systems*, volume 33, pp. 1877–1901, Virtual, 2020. Curran Associates, Inc.

534 Loïc Cordone, Benoît Miramond, and Philippe Thierion. Object detection with spiking neural net-  
535 works on automotive event data. In *International Joint Conference on Neural Networks*, pp. 1–8.  
536 IEEE, 2022.

538 Manon Dampfhofer and Thomas Mesquida. Neuromorphic lip-reading with signed spiking gated  
539 recurrent units. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern  
Recognition*, pp. 2141–2151, 2024.

540 Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha  
 541 Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, Yuyun Liao, Chit-Kwan Lin,  
 542 Andrew Lines, Ruokun Liu, Deepak Mathai, Steven McCoy, Arnab Paul, Jonathan Tse, Gu-  
 543 ruguhanathan Venkataraman, Yi-Hsin Weng, Andreas Wild, Yoonseok Yang, and Hong Wang.  
 544 *Loihi: a neuromorphic manycore processor with on-chip learning.* *IEEE Micro*, 38(1):82–99,  
 545 2018. doi: 10.1109/MM.2018.112130359.

546 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of  
 547 deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and  
 548 Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the*  
 549 *Association for Computational Linguistics*, pp. 4171–4186, Minneapolis, Minnesota, June 2019a.  
 550 Association for Computational Linguistics. doi: 10.18653/v1/N19-1423.

551 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep  
 552 bidirectional transformers for language understanding. In *Proceedings of the 2019 conference*  
 553 *of the North American chapter of the association for computational linguistics*, pp. 4171–4186,  
 554 2019b.

555 Kangrui Du, Yuhang Wu, Shikuang Deng, and Shi Gu. Temporal flexibility in spiking neural net-  
 556 works: Towards generalization across time steps and deployment friendliness. In *International*  
 557 *Conference on Learning Representations*, 2025.

558 Wei Fang, Zhaofei Yu, Yanqi Chen, Timothée Masquelier, Tiejun Huang, and Yonghong Tian. In-  
 559 corporating learnable membrane time constant to enhance learning of spiking neural networks.  
 560 In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2661–2671,  
 561 2021.

562 Wei Fang, Yanqi Chen, Jianhao Ding, Zhaofei Yu, Timothée Masquelier, Ding Chen, Liwei Huang,  
 563 Huihui Zhou, Guoqi Li, and Yonghong Tian. Spikingjelly: An open-source machine learning  
 564 infrastructure platform for spike-based intelligence. *Science Advances*, 9(40):eadi1480, 2023.

565 Matthias Fey and Jan E. Lenssen. Fast graph representation learning with PyTorch Geometric. In  
 566 *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.

567 Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi,  
 568 Stefan Leutenegger, Andrew J. Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza.  
 569 Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*,  
 570 44(1):154–180, 2022. doi: 10.1109/TPAMI.2020.3008413.

571 Daniel Gehrig, Antonio Loquercio, Konstantinos G Derpanis, and Davide Scaramuzza. End-to-end  
 572 learning of representations for asynchronous event-based data. In *Proceedings of the IEEE/CVF*  
 573 *International Conference on Computer Vision*, pp. 5633–5643, 2019.

574 Rafael C Gonzalez. *Digital image processing*. Pearson education india, India, 2009.

575 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad  
 576 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd  
 577 of models. *arXiv preprint arXiv:2407.21783*, 2024.

578 Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David  
 579 Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti  
 580 Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández  
 581 del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy,  
 582 Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array pro-  
 583 gramming with numpy. *Nature*, 585(7825):357–362, Sep 2020a. ISSN 1476-4687. doi:  
 584 10.1038/s41586-020-2649-2.

585 Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David  
 586 Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array program-  
 587 ming with numpy. *Nature*, 585(7825):357–362, 2020b.

588 Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with GPUs. *IEEE*  
 589 *Transactions on Big Data*, 7(3):535–547, 2019.

594     Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are  
 595     RNNs: Fast autoregressive transformers with linear attention. In Hal Daumé III and Aarti Singh  
 596     (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of  
 597     *Proceedings of Machine Learning Research*, pp. 5156–5165, Virtual, 13–18 Jul 2020. PMLR.

598     Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Fractalnet: Ultra-deep neural net-  
 599     works without residuals. *arXiv preprint arXiv:1605.07648*, 2016.

600     Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444,  
 601     May 2015. ISSN 1476-4687. doi: 10.1038/nature14539.

602     Gregor Lenz, Kenneth Chaney, Sumit Bam Shrestha, Omar Oubari, Serge Picaud, and Guido  
 603     Zarrella. Tonic: event-based datasets and transformations., July 2021. Documentation available  
 604     under <https://tonic.readthedocs.io>.

605     Patrick Lichtsteiner, Christoph Posch, and Tobi Delbrück. A  $128 \times 128$  120 db  $15 \mu\text{s}$  latency asyn-  
 606     chronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits*, 43(2):566–576,  
 607     2008. doi: 10.1109/JSSC.2007.914337.

608     XiuHong Lin, Changjie Qiu, Siqi Shen, Yu Zang, Weiquan Liu, Xuesheng Bian, Matthias Müller,  
 609     Cheng Wang, et al. E2pnet: event to point cloud registration with spatio-temporal representation  
 610     learning. *Advances in Neural Information Processing Systems*, 36:18076–18089, 2023.

611     Zhixuan Lin, Evgenii Nikishin, Xu He, and Aaron Courville. Forgetting transformer: Softmax  
 612     attention with a forget gate. In *International Conference on Learning Representations*, 2025.

613     Min Liu and Tobi Delbrück. Adaptive time-slice block-matching optical flow algorithm for dynamic  
 614     vision sensors. In *British Machine Vision Conference*, pp. 88. BMVA Press, 2018.

615     Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *Inter-  
 616     national Conference on Learning Representations*, 2017.

617     Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Confer-  
 618     ence on Learning Representations*, 2019.

619     Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models.  
 620     *Neural Networks*, 10(9):1659–1671, 1997.

621     Carver Mead. Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10):1629–1636,  
 622     1990. doi: 10.1109/5.58356.

623     Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, Andrew S Cassidy, Jun Sawada, Philipp  
 624     Akopyan, Bryan L Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, et al. A million spiking-  
 625     neuron integrated circuit with a scalable communication network and interface. *Science*, 345  
 626     (6197):668–673, 2014.

627     Nico Messikommer, Daniel Gehrig, Antonio Loquercio, and Davide Scaramuzza. Event-based asyn-  
 628     chronous sparse convolutional networks. 2020.

629     Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed represen-  
 630     tations of words and phrases and their compositionality. In C.J. Burges, L. Bottou, M. Welling,  
 631     Z. Ghahramani, and K.Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*,  
 632     volume 26, Lake Tahoe, Nevada, USA, 2013. Curran Associates, Inc.

633     Sharvil Nanavati. Haste: a fast, simple, and open rnn library. [https://github.com/  
 634     lmtn-com/haste/](https://github.com/lmtn-com/haste/), Jan 2020.

635     Hoang Nt and Takanori Maehara. Revisiting graph neural networks: All we have is low-pass filters.  
 636     *arXiv preprint arXiv:1905.09550*, 2019.

637     Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor  
 638     Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward  
 639     Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,  
 640     Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance  
 641     deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and  
 642     R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32, 2019.

648 Yansong Peng, Yueyi Zhang, Zhiwei Xiong, Xiaoyan Sun, and Feng Wu. Get: Group event trans-  
 649 former for event-based vision. In *Proceedings of the IEEE/CVF International Conference on*  
 650 *Computer Vision*, pp. 6038–6048, 2023.

651

652 Christoph Posch, Daniel Matolin, and Rainer Wohlgenannt. A qvga 143 db dynamic range frame-  
 653 free pwm image sensor with lossless pixel-level video compression and time-domain cds. *IEEE*  
 654 *Journal of Solid-State Circuits*, 46(1):259–275, 2011. doi: 10.1109/JSSC.2010.2085952.

655

656 Ofir Press, Noah A Smith, and Mike Lewis. Train short, test long: Attention with linear biases  
 657 enables input length extrapolation. *arXiv preprint arXiv:2108.12409*, 2021.

658

659 Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. Events-to-video: Bringing  
 660 modern computer vision to event cameras. In *Proceedings of the IEEE/CVF Conference on Com-*  
 661 *puter Vision and Pattern Recognition*, pp. 3857–3866, 2019.

662

663 Hongwei Ren, Yue Zhou, Jiadong Zhu, Xiaopeng Lin, Haotian Fu, Yulong Huang, Yuetong Fang,  
 664 Fei Ma, Hao Yu, and Bojun Cheng. Rethinking efficient and effective point-based networks for  
 665 event camera classification and regression. *IEEE Transactions on Pattern Analysis and Machine*  
 666 *Intelligence*, 47(8):6228–6241, 2025. doi: 10.1109/TPAMI.2025.3556561.

667

668 Kaushik Roy, Akhilesh Jaiswal, and Priyadarshini Panda. Towards spike-based machine intelligence  
 669 with neuromorphic computing. *Nature*, 575(7784):607–617, 2019.

670

671 Alberto Sabater, Luis Montesano, and Ana C Murillo. Event transformer+. a multi-purpose so-  
 672 lution for efficient event data processing. *IEEE transactions on Pattern Analysis and Machine*  
 673 *Intelligence*, 45(12):16013–16020, 2023.

674

675 Riccardo Santambrogio, Marco Cannici, and Matteo Matteucci. Farse-cnn: Fully asynchronous,  
 676 recurrent and sparse event-based cnn. In *European Conference on Computer Vision*, pp. 1–18.  
 677 Springer, 2024.

678

679 Simon Schaefer, Daniel Gehrig, and Davide Scaramuzza. Aegnn: Asynchronous event-based graph  
 680 neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
 681 12361–12371, 2022. doi: 10.1109/CVPR52688.2022.01205.

682

683 M. Schuster and K.K. Paliwal. Bidirectional recurrent neural networks. *IEEE Transactions on*  
 684 *Signal Processing*, 45(11):2673–2681, 1997. doi: 10.1109/78.650093.

685

686 Yusuke Sekikawa, Kosuke Hara, and Hideo Saito. Eventnet: Asynchronous recursive event process-  
 687 ing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,  
 688 pp. 3887–3896, 2019.

689

690 Amos Sironi, Manuele Brambilla, Nicolas Bourdis, Xavier Lagorce, and Ryad Benosman. Hats:  
 691 Histograms of averaged time surfaces for robust event-based object classification. In *Proceedings*  
 692 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1731–1740, 2018.

693

694 Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: En-  
 695 hanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

696

697 Anand Subramoney, Khaleelulla Khan Nazeer, Mark Schöne, Christian Mayr, and David Kappel.  
 698 Efficient recurrent architectures through activity sparsity and sparse back-propagation through  
 699 time. In *International Conference on Learning Representations*, 2023.

700

701 Jingkai Sun, Qiang Zhang, Jiaxu Wang, Jiahang Cao, Hao Cheng, and Renjing Xu. Event  
 702 masked autoencoder: Point-wise action recognition with event-based cameras. In *IEEE In-*  
 703 *ternational Conference on Acoustics, Speech and Signal Processing*, pp. 1–5, 2025. doi:  
 704 10.1109/ICASSP49660.2025.10888760.

705

706 Ganchao Tan, Yang Wang, Han Han, Yang Cao, Feng Wu, and Zheng-Jun Zha. Multi-grained spatio-  
 707 temporal features perceived network for event-based lip-reading. In *Proceedings of the IEEE/CVF*  
 708 *Conference on Computer Vision and Pattern Recognition*, pp. 20094–20103, June 2022.

702 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
 703 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg,  
 704 S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural*  
 705 *Information Processing Systems*, volume 30, Long Beach, California, USA, 2017. Curran Asso-  
 706 *ciates, Inc.*

707 Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on*  
 708 *computer vision*, pp. 3–19, 2018.

709 Jiancheng Yang, Qiang Zhang, Bingbing Ni, Linguo Li, Jinxian Liu, Mengdie Zhou, and Qi Tian.  
 710 Modeling point clouds with self-attention and gumbel subset sampling. In *Proceedings of the*  
 711 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3323–3332, 2019.

712 Songlin Yang and Yu Zhang. Fla: A triton-based library for hardware-efficient implementa-  
 713 tions of linear attention mechanism, January 2024. URL [https://github.com/fla-org/](https://github.com/fla-org/flash-linear-attention)  
 714 [flash-linear-attention](https://github.com/fla-org/flash-linear-attention).

715 Man Yao, Guangshe Zhao, Hengyu Zhang, Yifan Hu, Lei Deng, Yonghong Tian, Bo Xu, and Guoqi  
 716 Li. Attention spiking neural networks. *IEEE transactions on Pattern Analysis and Machine*  
 717 *Intelligence*, 45(8):9393–9410, 2023.

718 Man Yao, Ole Richter, Guangshe Zhao, Ning Qiao, Yannan Xing, Dingheng Wang, Tianxiang Hu,  
 719 Wei Fang, Tugba Demirci, Michele De Marchi, Lei Deng, Tianyi Yan, Carsten Nielsen, Sadique  
 720 Sheik, Chenxi Wu, Yonghong Tian, Bo Xu, and Guoqi Li. Spike-based dynamic computing with  
 721 asynchronous sensing-computing neuromorphic chip. *Nature Communications*, 15(1):4464, May  
 722 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-47811-6.

723 Chengguo Yuan, Yu Jin, Zongzhen Wu, Fanting Wei, Yangzirui Wang, Lan Chen, and Xiao Wang.  
 724 Learning bottleneck transformer for event image-voxel feature fusion based classification. In  
 725 *Chinese Conference on Pattern Recognition and Computer Vision*, pp. 3–15. Springer, 2023.

726 Chenlin Zhou, Han Zhang, Liutao Yu, Yumin Ye, Zhaokun Zhou, Liwei Huang, Zhengyu Ma, Xi-  
 727 aopeng Fan, Huihui Zhou, and Yonghong Tian. Direct training high-performance deep spiking  
 728 neural networks: a review of theories and methods. *Frontiers in Neuroscience*, Volume 18 - 2024,  
 729 2024. ISSN 1662-453X. doi: 10.3389/fnins.2024.1383844.

730 Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang,  
 731 Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applica-  
 732 tions. *AI Open*, 1:57–81, 2020. ISSN 2666-6510. doi: <https://doi.org/10.1016/j.aiopen.2021.01.001>.

733 Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Unsupervised event-based  
 734 learning of optical flow, depth, and egomotion. In *Proceedings of the IEEE/CVF Conference on*  
 735 *Computer Vision and Pattern Recognition*, pp. 989–997, 2019.

736

## 737 A APPENDIX

### 738 A.1 DATASETS

739 The DVS Gesture dataset is a benchmark commonly employed for model evaluation. It comprises  
 740 11 categories of hand gestures and is officially divided into a training set with 1,176 samples and a  
 741 testing set with 288 samples.

742 The ASL-DVS dataset contains 24 classes corresponding to letters from American Sign Language,  
 743 amounting to a total of 100,800 samples. Each class consists of 4,200 samples, and the duration of  
 744 each sample is approximately 100 ms. Following the methodology in Bi et al. (2019), we partition  
 745 the dataset by allocating the initial 80% for training and the remaining 20% for testing.

746 The DVS-Lip dataset encompasses 100 word classes derived from the visual information of a  
 747 speaker’s lip movements. It provides an official training/testing split with 14,896 and 4,975 samples,  
 748 respectively.

756 Table 5: Hyper-parameters of training models for classification tasks on different datasets  
757

758      Dataset	$D$	$D_f$	$n_{head}$	Depth	Repeats	$n_{gpus}$	$lr_{min}$
759      DVS Gesture	64	128	2	4	24	4	0
760      ASL-DVS	64	128	2	2	1	7	$10^{-6}$
761      DVS-Lip	192	384	6	16	3	4	$10^{-6}$

762  
763 A.2 BI-DIRECTIONAL LINEAR ATTENTIONS  
764765 Ignoring the layer index for simplicity, the recurrence relation for the hidden state  $\mathbf{S}[t]$  at time-step  $t$   
766 in a standard linear attention model is:  
767

768 
$$\mathbf{S}[t] = f(\mathbf{S}[t-1], \mathbf{k}[t], \mathbf{v}[t]), \quad (8)$$

769 
$$\mathbf{O}[t] = \mathbf{W}_o \mathbf{S}[t], \quad (9)$$

770 where  $\mathbf{k}[t]$  and  $\mathbf{v}[t]$  are input-dependent key and value vectors,  $\mathbf{W}_o$  is the output projection matrix,  
771 and  $\mathbf{O}[t]$  is the output vector. The specific recurrence function,  $f$ , used in this work is based on the  
772 Forgetting Transformer (FOX) (Lin et al., 2025) implemented by the Flash Linear Attention library  
773 (Yang & Zhang, 2024)774 We adapt this formulation to be bi-directional by maintaining separate forward and backward hidden  
775 states, which are then concatenated and fused through a linear layer to produce the output  $\mathbf{O}_{fb}[t]$ :  
776

777 
$$\mathbf{S}_f[t] = f(\mathbf{S}_f[t-1], \mathbf{k}[t], \mathbf{v}[t]), \quad (10)$$

779 
$$\mathbf{S}_b[t] = f(\mathbf{S}_b[t-1], \mathbf{k}[L-t-1], \mathbf{v}[L-t-1]), \quad (11)$$

780 
$$\mathbf{S}_{fb}[t] = [\mathbf{S}_f[t]; \mathbf{S}_b[t]], \quad (12)$$

781 
$$\mathbf{O}_{fb}[t] = \mathbf{W}_{fb} \mathbf{S}_{fb}[t]. \quad (13)$$

782 Unlike classic bi-directional RNNs (Schuster & Paliwal, 1997) that often use independent parameters  
783 for each direction, our model employs shared parameters for the forward and backward passes.  
784 Consequently, the only increase in parameters compared to the uni-directional model arises from the  
785 output projection matrix  $\mathbf{W}_{fb}$ , which has twice the number of parameters as the original  $\mathbf{W}_o$ .  
786787 A.3 MODEL HYPER-PARAMETERS  
788789 Unless otherwise stated, all models were trained using BFloat16 mixed precision. The training  
790 configuration for all models includes a base learning rate of  $lr_b = 0.001$ , a batch size of 64, and the  
791 AdamW optimizer (Loshchilov & Hutter, 2019), conducted over 64 epochs. The effective learning  
792 rate is determined by a linear scaling rule based on the number of GPUs ( $n_{gpus}$ ) used in distributed  
793 data-parallel training:  $lr = lr_b \cdot n_{gpus} / 256$ . A warmup phase is implemented for the first four  
794 epochs, during which the learning rate is linearly increased from  $0.01 \cdot lr$  to its full value,  $lr$ . For  
795 the subsequent epochs, a cosine annealing schedule (Loshchilov & Hutter, 2017) is employed to  
796 gradually reduce the learning rate to a minimum value,  $lr_{min}$ . For the DVS Gesture and ASL-  
797 DVS datasets, both weight decay and label smoothing were disabled. In contrast, for the DVS-Lip  
798 classification task, we set the weight decay to 0.05 and applied label smoothing with a factor of 0.1.799 Table 5 provides a detailed summary of the model-specific hyper-parameters. Here,  $D$  denotes  
800 the embedding dimension,  $D_f$  represents the hidden feature dimension of the feed-forward neural  
801 network (FFN), and  $n_{head}$  is the total number of attention heads. The `repeats` parameter specifies  
802 how many times the training set is iterated through within a single epoch. Notably, the number of  
803 heads for the key ( $\mathbf{k}$ ) and value ( $\mathbf{v}$ ) projections is set to  $n_{head}/2$ , and group normalization (Wu &  
804 He, 2018) is applied to both. To prevent exploding gradients, we employ gradient clipping, capping  
805 the  $L_2$  norm of the gradients at 1.0.806 For the DVS Gesture classification model, the output of each FFN is average-pooled with a window  
807 size of 2, whereas other models do not use pooling. The model for the DVS-Lip classification task  
808 was pre-trained on the DVS-Lip dataset using a self-supervised learning approach. This pre-training  
809 phase utilized a minimum learning rate of  $lr_{min} = 10^{-6}$ , a weight decay of 0.05, a `repeats` value  
of 3, and a masking ratio of 30%. Refer to Appendix A.5 for more details.

810  
811 A.4 DATA AUGMENTATIONS812 Denote  $\mathcal{U}(a, b)$  as the uniform distribution between  $a$  and  $b$ , and  $\text{RandInt}(m, n)$  as a random integer  
813 taken from the set  $m, m + 1, \dots, n$ , where each integer has an equal probability of being selected.814 For an event stream, the data augmentations are applied on events directly. For simplicity, we omit  
815 the event index in this subsection. Unless otherwise specified, augmentations are applied on each  
816 event stream independently. Note that the coordinates are converted to floating precision before  
817 applying any augmentation. After all augmentations are applied, coordinates will be quantized, and  
818 only events whose coordinates are valid, i.e.,  $x \in [0, W - 1], y \in [0, H - 1]$ , are kept.819 For the classification task on DVS Gesture, the following transformations are each applied indepen-  
820 dently with a probability of 0.6:  
821

- **Random Resizing:** Coordinates  $(x, y)$  are scaled to  $(s_x \cdot x, s_y \cdot y)$ , with scaling factors  $s_x, s_y \sim \mathcal{U}(0.8, 1.2)$ .
- **Random Rotation:** Coordinates are rotated by an angle  $r \sim \mathcal{U}(-10, 10)$  degrees.
- **Random Shearing:** A shear transformation is applied with factors  $\lambda_x, \lambda_y \sim \mathcal{U}(-0.02, 0.02)$ .
- **Random Translation:** Coordinates are translated by offsets  $d_x, d_y \sim \mathcal{U}(-16, 16)$ .
- **Random Erasing:** Erase an  $h \times w$  area with  $h, w \sim \mathcal{U}(0, 16)$  with the probability 0.1. The center of this area  $(c_x, c_y)$  satisfy  $c_x \sim \mathcal{U}(0, W - 1), c_y \sim \mathcal{U}(0, H - 1)$ .
- **Temporal Chunk Dropout:** A number of temporal chunks,  $n_r = \text{RandInt}(0, 8)$ , are removed from the event stream. The length of each removed chunk,  $l_{\text{chunk}}$ , is determined relative to the total stream length,  $L$ , according to the sampling distribution  $l_{\text{chunk}} = \frac{\text{RandInt}(1, 256)}{L}$ .

835 No data augmentations were applied for the ASL-DVS dataset.  
836837 During the self-supervised phase of the model for classifying DVS-Lip, a series of geometric trans-  
838 formations are employed. Each of the following augmentations is applied independently with a  
839 probability of 0.5:

- **Random Resizing:** Coordinates  $(x, y)$  are scaled to  $(s_x \cdot x, s_y \cdot y)$ , with scaling factors  $s_x, s_y \sim \mathcal{U}(0.8, 1.2)$ .
- **Random Rotation:** Coordinates are rotated by an angle  $r \sim \mathcal{U}(-15, 15)$  degrees.
- **Random Shearing:** A shear transformation is applied with factors  $\lambda_x, \lambda_y \sim \mathcal{U}(-0.05, 0.05)$ .
- **Horizontal Flipping:** The event stream is flipped horizontally.
- **Random Translation:** Coordinates are translated by offsets  $d_x, d_y \sim \mathcal{U}(-16, 16)$ .

840 When training the model for classifying DVS-Lip, we use the following data augmentations:  
841

- **Random Resizing:** Resize  $(x, y)$  to  $(s_x \cdot x, s_y \cdot y)$  where  $s_x, s_y \sim \mathcal{U}(0.8, 1.2)$ .
- **Random Rotation:** Coordinates are rotated by an angle  $r \sim \mathcal{U}(-15, 15)$  degrees.
- **Random Shearing:** Shear transform on  $x$  and  $y$  with shear factors  $\lambda_x, \lambda_y \sim \mathcal{U}(-0.05, 0.05)$ .
- **Random Flip:** The event stream is flipped horizontally with a probability of 0.5.
- **Random Translation:** Translate  $x$  and  $y$  with translations  $d_x, d_y \sim \mathcal{U}(-16, 16)$ .
- **Random Erasing:** Erase an  $h \times w$  area with  $h, w \sim \mathcal{U}(0, 16)$  with the probability 0.1. The center of this area  $(c_x, c_y)$  satisfies  $c_x \sim \mathcal{U}(0, W - 1), c_y \sim \mathcal{U}(0, H - 1)$ .

860 The augmentations listed above are each applied independently with a probability of 0.5. The token-  
861 mix is applied on the embedding tensor with probability 0.5. Specifically, when training on cluster  
862 events, the intensity  $\rho$  is randomly set to 1 with a probability of 0.1. We use drop path (Larsson  
863 et al., 2016) in the linear attention layer, with the probability increasing linearly from 0 to 0.4 with  
depth.

864  
865

## A.5 SELF-SUPERVISED TRAINING DETAILS

866  
867  
868  
869  
870  
871

The event-wise nature of the event2vec representation lends itself well to self-supervised pre-training, which can significantly enhance model performance. Specifically, we adopt a masked modeling approach, akin to that used in BERT. The training objective is to mask the spatial coordinates  $(x, y, p)$  of a subset of these events and train the model to predict the masked coordinates based on the context provided by the surrounding events and their associated temporal information. This task compels the model to learn a meaningful understanding of spatio-temporal event patterns.

872  
873  
874  
875

The self-supervised training framework is analogous to the Masked Language Model (MLM) objective in BERT (Devlin et al., 2019a). Given a batch of embedding tensors  $\mathbf{v}$  of shape  $(B, L, D)$ , where  $B$  is the batch size,  $L$  is the sequence length, and  $D$  is the embedding dimension, the process begins by randomly masking a portion of the input tokens.

876  
877  
878  
879  
880  
881

A binary mask  $\mathbf{m}$  of shape  $(B, L)$  is generated from a Bernoulli distribution. The probability of masking any given token is set to 0.3, which defines the mask ratio. Each token  $\mathbf{v}[i][j]$  corresponding to a mask entry  $\mathbf{m}[i][j] = 1$  is replaced by a single, learnable,  $D$ -dimensional mask token  $\mathbf{v}_m$ . This operation results in a corrupted embedding tensor, denoted as  $\hat{\mathbf{v}}$ . Concurrently, the original coordinates  $(\mathbf{x}_m, \mathbf{y}_m, \mathbf{p}_m)$  of the masked tokens are preserved to serve as the ground truth for the reconstruction loss.

882  
883  
884

The corrupted tensor  $\hat{\mathbf{v}}$  is then processed by the model’s linear attention layers. Following this, the output embeddings that correspond to the initially masked positions, denoted  $\hat{\mathbf{v}}_m$ , are extracted from the final output tensor using the mask  $\mathbf{m}$ .

885  
886  
887  
888

The objective is for the model to reconstruct the original spatial and polarity information from these corrupted embeddings. To achieve this, we first apply the inverse of the spatio-temporal fusion operation to isolate the spatial component of the reconstructed embeddings:

889  
890

$$\hat{\mathbf{v}}_s = \frac{\hat{\mathbf{v}}_m}{\rho} - \mathbf{v}_t. \quad (14)$$

891  
892  
893  
894  
895  
896  
897

The resulting tensor,  $\hat{\mathbf{v}}_s$ , is treated as the reconstructed spatial embedding. It is then passed through a decoder network, which mirrors the architecture of the spatial embedding encoder, to predict the original coordinates  $(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{p}})$ . Specifically, this decoder consists of a stack of linear layers, Layer Normalization, and ReLU activation functions. The network is designed to gradually reduce the feature dimension from  $D$  down to 3. The final output layer uses a tanh activation function to constrain the predicted values to the range  $(-1, 1)$ . This aligns with the input preprocessing, where the ground-truth coordinates are also normalized to the same range.

898  
899

Finally, the training objective is to minimize the Mean Squared Error (MSE) loss between the predicted coordinates  $(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{p}})$  and the ground-truth coordinates  $(\mathbf{x}_m, \mathbf{y}_m, \mathbf{p}_m)$  of the masked tokens.

900  
901

## A.6 EXPERIMENTAL DETAILS FOR PRE-PROCESSING LATENCY

902  
903  
904  
905  
906  
907  
908  
909  
910  
911

The pre-processing latency benchmarks, with results presented in Table 2, were conducted on a Red Hat Enterprise Linux 8.10 server. This server was equipped with an NVIDIA A100 GPU (80GB PCIe), an Intel Xeon Gold 6326 CPU (utilizing 8 cores), and 256GB of RAM. To mitigate the impact of data I/O, the DVS Gesture dataset was loaded entirely into RAM for the duration of the experiments. For each method, we optimized the batch size and the number of workers to achieve the minimum possible latency. All implementations were based on the publicly available source code from their respective original publications or other high-performance software libraries. The reported latency for each method is the average of 8 measurement runs, which were preceded by 2 warm-up runs to ensure system stability. Further implementation details are provided in Table 6.

912  
913

## A.7 EXPERIMENTAL DETAILS FOR MODEL THROUGHPUT AND LATENCY

914  
915  
916

The throughput and latency experiments, with results presented in Table 3, were conducted under the same operating system and hardware environment as the pre-processing benchmarks, which are detailed in Appendix A.6. The dataset was also pre-loaded into RAM to eliminate I/O bottlenecks.

917

For each model, we performed a search for the optimal batch size to maximize performance. The optimal batch sizes were determined to be 64 for Sparse GRU, 16 for FARSE-CNN, 64 for AEGNN,

918  
919  
920 Table 6: Implementation details of Table 2  
921  
922  
923  
924  
925  
926

Method	Batch size	Workers	Software Library
Random Sampling	64	64	Numpy (Harris et al., 2020b)
Frame	8	8	SpikingJelly (Fang et al., 2023)
Graph	8	8	PyTorch Geometric (Fey & Lenssen, 2019)
Voxel Grid	16	8	Tonic (Lenz et al., 2021)
K-Means	2	2	Faiss (Johnson et al., 2019)
Window Slicing	8	8	Original paper (Santambrogio et al., 2024)

927  
928 64 for SNN and 512 for event2vec, respectively. We observed that a fixed number of 8 workers  
929 yielded the best performance across all models. The benchmarking process for each model involved  
930 an initial warm-up phase, followed by multiple measurement runs, the results of which were then  
931 averaged. Due to significant variations in computational cost among the models, the number of  
932 warm-up iterations and measurement batches was tailored for each specific model. However, we  
933 ensured that the number of batches was sufficiently large, such that further increases did not yield  
934 any significant changes in the measured performance, confirming the stability of our results.

935 Regarding the implementations, for Sparse GRU<sup>1</sup> and FARSE-CNN<sup>2</sup>, the complete source code  
936 was available in the official GitHub repositories, and we used them directly. For AEGNN, only  
937 the inference code was publicly available<sup>3</sup>. We therefore implemented the necessary Dataset  
938 class and model architecture by referencing their provided implementation for the N-Cars dataset  
939 (Sironi et al., 2018). For SNNs, the model used by Yao et al. (2023) enhances the Parametric Leaky  
940 Integrate-and-Fire Neuron Network (PLIF-Net) (Fang et al., 2021) with attentions, but the source  
941 code is not released. Given these additional attention modules only add slight complexity, we evaluate  
942 on the PLIF-Net as an alternative. SpikingJelly (Fang et al., 2023) provides a high-performance  
943 implementation<sup>4</sup> for the PLIF-Net with advanced accelerating techniques, and we benchmark on  
944 code from SpikingJelly directly.

945 All benchmarks were running in BFloat16 mixed precision except for Sparse GRU, which depends  
946 on the Haste library (Nanavati, 2020) with only supports Float32.

#### 947 948 A.8 EXPERIMENTAL DETAILS FOR THE ABLATION STUDY ON EVENT NUMBERS

949 The ablation study on the number of events, with results reported in Figure 3, evaluates the impact  
950 on model throughput, latency, and accuracy. These experiments were conducted using the same  
951 operating system and hardware environment detailed in Appendix A.6. During the training process  
952 on the DVS Gesture dataset, the chunk length parameter  $l_{chunk}$  for the temporal chunk dropout  
953 augmentation is scaled proportionally with the number of events. Furthermore, for experiments  
954 using 64 and 32 events, all data augmentation techniques are disabled. This measure is implemented  
955 to prevent the augmentation from inadvertently removing all events, which would subsequently lead  
956 to a NaN (Not-a-Number) loss.

#### 957 958 A.9 VISUALIZATION OF NEIGHBORHOOD SEMANTICS

959 Due to space constraints in the main paper, Figures 5(a) and 5(c) display visualizations for only  
960 a single event polarity. For completeness, this section provides supplementary visualizations that  
961 include both polarities. Figure 7 illustrates the embedding weights mapped to the RGB color space,  
962 while Figure 8 depicts them as a vector field.

963  
964  
965  
966  
967  
968 <sup>1</sup><https://github.com/Efficient-Scalable-Machine-Learning/EvNN>

969 <sup>2</sup><https://github.com/AIRLab-POLIMI/farse-cnn>

970 <sup>3</sup><https://github.com/uzh-rpg/aegnn>

971 <sup>4</sup>[https://spikingjelly.readthedocs.io/zh-cn/latest/activation\\_based\\_en/classify\\_dvsg.html](https://spikingjelly.readthedocs.io/zh-cn/latest/activation_based_en/classify_dvsg.html)

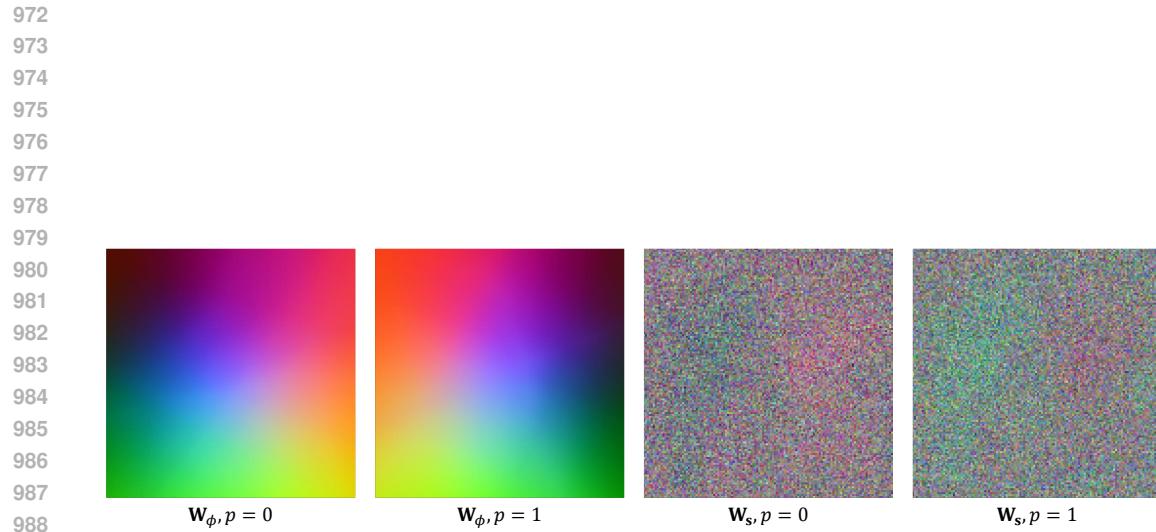


Figure 7: Visualization of the parametric embedding weight  $\mathbf{W}_\phi$  and the standard embedding weight  $\mathbf{W}_s$  in the RGB domain.

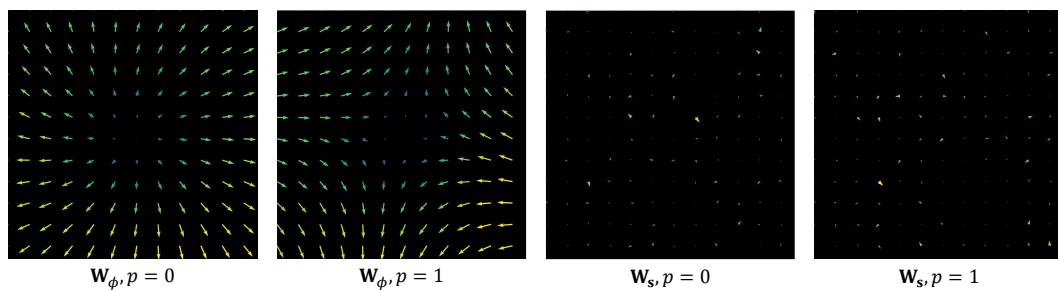


Figure 8: Visualization of the parametric embedding weight  $\mathbf{W}_\phi$  and the standard embedding weight  $\mathbf{W}_s$  in the vector field.