

OmniEvent: Unified Event Representation Learning

WeiQi Yan^{1*}, Chenlu Lin^{1*}, Youbiao Wang¹, Zhipeng Cai^{2†}, Xiuhong Lin¹, Yangyang Shi²,
Weiquan Liu³, Yu Zang^{1†}

¹Fujian Key Lab of Sensing and Computing for Smart Cities, School of Informatics, Xiamen University, Xiamen, China

²Meta, Menlo Park, USA

³College of Computer Engineering, Jimei University, Xiamen, China

{yanweiqi199888, czptc2h, shiyang1983}@gmail.com, {36920231153214, 36920241153251, xhlinxm}@stu.xmu.edu.cn, wqliu@jmu.edu.cn, zangyu7@126.com

Abstract

Event cameras have gained increasing popularity in computer vision due to their ultra-high dynamic range and temporal resolution. However, event networks heavily rely on task-specific designs due to the unstructured data distribution and spatial-temporal (S-T) inhomogeneity, making it hard to reuse existing architectures for new tasks. We propose OmniEvent, an innovative unified event representation learning framework that achieves SOTA performance across diverse tasks, fully removing the need for task-specific designs. Unlike previous methods that treat event data as 3D point clouds with manually tuned S-T scaling weights, OmniEvent proposes a decouple-enhance-fuse paradigm, where the local feature aggregation and enhancement are done independently on the spatial and temporal domains to avoid inhomogeneity issues. Space-filling curves are applied to enable large receptive fields while improving memory and compute efficiency. The features from individual domains are then fused by attention to learn S-T interactions. The output of OmniEvent is a grid-shaped tensor, which enables standard vision models to process event data without architectural changes. With a unified framework and similar hyperparameters, OmniEvent outperforms (task-specific) SOTA by up to 68.2% across 3 representative tasks and 10 datasets.

Code — <https://github.com/WickyAn/OmniEvent>

Extended version — <https://arxiv.org/abs/2508.01842>

1 Introduction

Event cameras are a new type of camera that only respond to pixel positions where the illumination change exceeds a certain threshold. Compared with conventional cameras, event cameras have very high dynamic range (HDR, up to 120 dB (Gallego et al. 2020)) and temporal resolution, very low response latency and power consumption. These advantages make event cameras less susceptible to challenging illumination and motion. However, the output of an event camera consists of a stream of events that only encode the time, location and polarity (direction) of brightness changes. Consequently, each event alone carries very little information, and

*Equal contribution.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

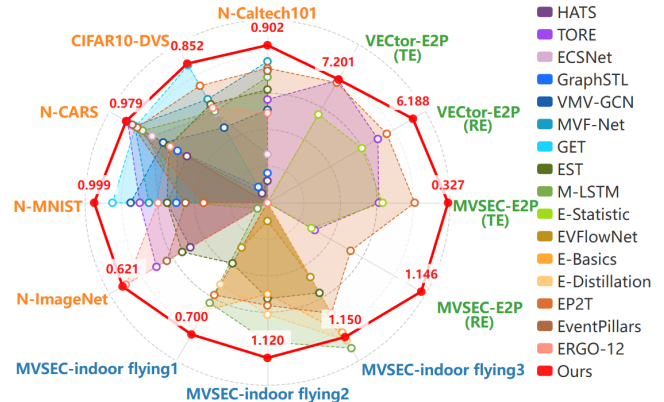


Figure 1: **Teaser.** We propose OmniEvent, a unified event representation framework that achieves SOTA performance across diverse tasks (Object Classification, 2D-3D Registration and Optical Flow Estimation, the color of each dataset represents the task) without complex task-specific design or hyperparameter tuning. (RE) and (TE) are rotation and translation errors.

the event generation rate varies significantly in different scenarios. Furthermore, the unstructured data stream of event cameras is sparse in the spatial plane but may be dense (up to 300 million events per second (Son et al. 2017)) in the temporal dimension. All of these challenges have spurred the community to develop specialized algorithms tailored to the unique data characteristics of event cameras.

Data-driven approaches have gained significant attention in recent years. Given that event cameras are fundamentally 2D visual sensors, an intuitive approach is to leverage existing research on frame-based convolutional neural networks (CNNs). To accomplish this, asynchronous event data is typically transformed into a grid-shaped feature tensor representation that can be synchronously updated. Existing approaches predominantly rely on hand-crafted event tensorization methods, such as statistical features (Maqueda et al. 2018; Nguyen et al. 2019; Park et al. 2016; Zhu et al. 2018b), exponential (Lagorce et al. 2016; Sironi et al. 2018) or linear (Cannici et al. 2019) decays, event surfaces (Almtrafi et al. 2020; Baldwin et al. 2022), and event voxel (Zhu et al. 2019; Rebecq et al. 2019). Constrained by their repre-

sentational capacity, these hand-crafted event representation methods often lose S-T details and cannot be optimized for different downstream tasks.

Recent approaches in event-based representation learning have emerged by treating event streams as 3D point clouds. However, classical frameworks (Qi et al. 2017b; Sekikawa, Hara, and Saito 2019) face dual challenges when applied to event data. First, Euclidean distance-based neighborhood construction struggles to adapt to spatiotemporal heterogeneity—fixed spatial-radius neighborhood searches prove inefficient for dynamically sparse event distributions, while microsecond-level timestamp differences are nearly obliterated in Euclidean distance calculations. Second, although the K-Nearest Neighbors (KNN) algorithm enables local feature aggregation, its receptive field remains constrained by local neighborhood size. Increasing the neighbor count K expands coverage but incurs heavy compute and memory costs, failing to fundamentally resolve distance degradation in high-dimensional spaces. Existing improvements, such as EP2T (Lin et al. 2024), attempt to mitigate these issues through spatiotemporal feature-weighted fusion. However, static weighting mechanisms cannot adapt to dynamic requirements across diverse scenarios.

We propose *OmniEvent*, an innovative unified representation learning framework effective for different event tasks. *OmniEvent* first decouples event data into spatial and temporal feature subspaces, where independent encoders capture high-precision geometric features and microsecond-level motion patterns. Subsequently, space-filling curve-guided feature aggregation maps 3D spatial coordinates into 1D continuous curve sequences. This strategy naturally establishes long-range spatial correlations through curve traversal, avoiding KNN computational bottlenecks while enabling the model to adaptively adjust fusion weights between local and global features. Combining with standard architectures (He et al. 2016; Zhu et al. 2018b; Pham et al. 2020), *OmniEvent* surpasses state-of-the-art methods across 3 event tasks and 10 datasets, advancing SOTA by a large margin (up to 68% error reduction). Our main *contributions* are summarized below:

- We propose *OmniEvent*, an innovative framework that can learn effective representations for various event vision tasks with a unified architecture.
- We introduce the decouple-enhance-fuse paradigm, which *fully* decouples spatial and temporal domains during early feature aggregation, eliminating their interference without task/dataset-specific scaling weight tuning.
- We apply space-filling curves for event representation learning which significantly enlarge the receptive fields and removes the compute and memory bottlenecks from KNN-based methods.

2 Related Work

Hand-crafted Representations. Direct projections mark event pixel positions but ignore temporal information (Maqueda et al. 2018). Statistical features (event counts/timestamps) (Zhu et al. 2018b; Nguyen et al. 2019; Park et al. 2016) and voxels (Zhu et al. 2019; Rebecq et al.

2019) approximate spatiotemporal distributions but lack fine-grained correlations. Time decay representations (exponential/linear) (Lagorce et al. 2016; Cannici et al. 2019) and event surfaces (Benosman et al. 2013; Almatrafi et al. 2020) preserve details through SNN-inspired structures (Tavanaei et al. 2019), with improved versions using event buffering (Baldwin et al. 2022) or neighborhood merging (Sironi et al. 2018; Manderscheid et al. 2019) for stability. Graph-based methods (Zhou et al. 2021; Mitrokhin et al. 2020; Schaefer, Gehrig, and Scaramuzza 2022) excel spatially but lose temporal details with high computational costs. Overall, these representations fundamentally limit comprehensive spatiotemporal capture, hindering task adaptability.

Learned Representations. SNNs directly process events but require specialized hardware (Paredes-Vallés, Scheper, and De Croon 2019; Andreopoulos et al. 2018) and face training challenges (Huh and Sejnowski 2018; Rueckauer et al. 2017). EST/DDES (Gehrig et al. 2019; Tulyakov et al. 2019) transform events via MLP-learned temporal kernels but remain representationally limited. Matrix-LSTM (Cannici et al. 2020) encodes patch sequences but handles only short-range dependencies with custom parallelization needs. ECSNet (Chen et al. 2022b) and EP2T process events as point clouds but suffer from Euclidean metric constraints and inefficient long-range modeling. *OmniEvent* overcomes these by spatiotemporally decoupling dimensions, improving efficiency and correlation capture.

3 Method

During event data collection, each pixel of the event camera responds independently and asynchronously to illumination changes. Specifically, given an event camera of size (H, W) , and $I(t_i, h, w)$ as the illumination at time t_i for the pixel at (h, w) . When the inequality

$$\|\log(I(t_i, h, w)) - \log(I(t_{i-1}, h, w))\| > \tau \quad (1)$$

is satisfied, this pixel will generate an *event*, which is expressed as $\mathbf{e} = (h, w, t_i, p)$. The polarity $p \in \{+1, -1\}$ indicates the brightness change direction (increasing or decreasing) and τ is the trigger threshold. t_{i-1} represents the time when the event was triggered last time at pixel (h, w) .

OmniEvent treats event data as spatio-temporal point clouds. As the imaging plane (H, W) and the time axis can be combined to establish a three-dimensional spatio-temporal coordinate system (h, w, t_i) . Such a representation preserves the original information and data sparsity.

As illustrated in Fig. 2, *OmniEvent* consists of five consecutive modules: Event Fusion and Sampling (EFS), Spatio-Temporal Decoupling (STD), Space Filling Curve based Aggregation (SFCA), Spatio-temporal Separated Attention (STA), and Feature Tensorization (FT).

The framework design follows a *decouple-enhance-fuse* paradigm. **EFS** (Sec. 3.1) fuses the event points and subsamples them into a fixed number to enable efficient computation and minimize information loss. **STD** (Sec. 3.2) proposes a radical decoupling of spatial and temporal dimensions. This approach reduces the dependence on joint distance metric optimization and enables the capture of long-

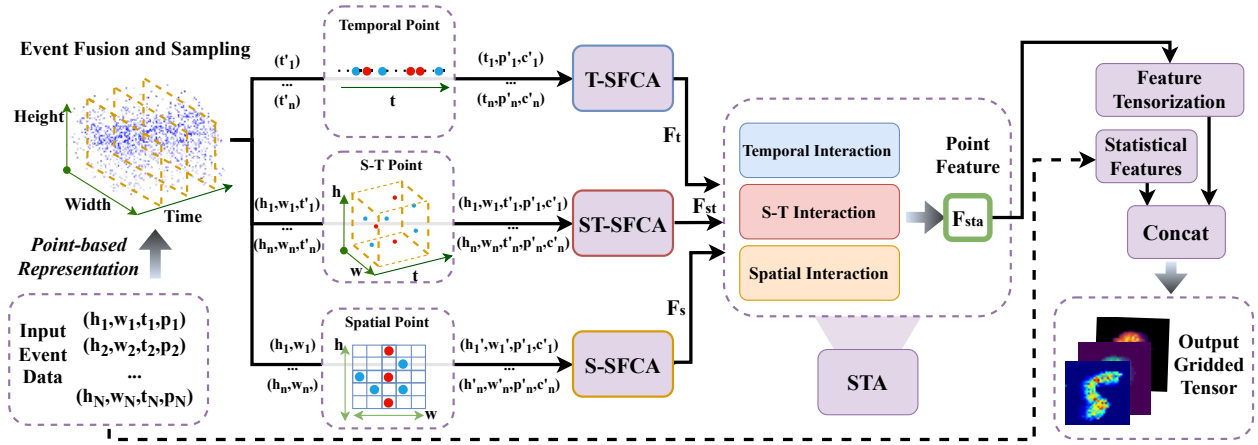


Figure 2: **OmniEvent architecture.** OmniEvent treats event data as spatio-temporal point clouds. Event Fusion and Sampling (EFS) fuses and samples unstructured event data into a fixed number of points. Three sets of Space-Filling-Curve-based Aggregation (SFCA) modules independently extract spatial-only, temporal-only and spatial-temporal information using different distance metrics and features, effectively addressing spatio-temporal inhomogeneity. Subsequently, Spatio-temporal Separated Attention (STA) models the spatio-temporal correlations of features using attention. Finally, Feature Tensorization (FT) converts unstructured event features into a grid-shaped tensor and combines them with event statistical features. The output grid-shaped tensor can be used by any frame-based method.

range spatiotemporal correlations through separate temporal and spatial neighborhood definitions. **SFCA** (Sec. 3.3) conducts sequential feature enhancement on decoupled features by fully utilizing event presentation. SFCA also decouples positional embeddings from feature aggregation and the event fusion sampling method, achieves SOTA representation capabilities while improving algorithm efficiency. **STA** (Sec. 3.4) reorganizes the decoupled temporal and spatial features into global features, modeling spatiotemporal correlations. **FT** (Sec. 3.4) uses sparse convolution to convert unstructured spatiotemporal point features into grid-shaped tensors and introduces event statistical features to compensate for the global information loss due to sampling. FT enables OmniEvent to focus on extracting spatiotemporal correlations and detailed information. The output grid-shaped tensor can be integrated into any frame-based network, allowing it to process event data without changing the architecture/training process.

3.1 Event Fusion and Sampling (EFS)

Although 3D point-based architectures (Qi et al. 2017b; Li et al. 2018; Liu et al. 2020) can adapt to irregular and discrete event point clouds, they are still limited by the number of points they can process simultaneously. Fortunately, we found significant information redundancy within events. For instance, moving objects with high contrast against the background may trigger numerous events simultaneously at the same pixel location (see Eq.(1)). This redundancy allows us to perform event point fusion before sampling.

As illustrated in the top left of Fig. 2, similar to event voxelization (Zhu et al. 2019; Rebecq et al. 2019), we divide a batch of events into T temporal segments, where events within the same pixel and temporal segment are fused. Specifically, given N events $\mathbf{e}_{hw} = (h, w, t_i, p_i)$,

$i \in \{1, \dots, c\}$ on position (x, y) in time slice T_n , the fused event can be represented as $\mathbf{e}_{hw}^l = (h, w, t_{avg}, p_{acc}, c)$, $t_{avg} = \frac{1}{c} \sum_i t_i$ and $p_{acc} = \sum_i p_i$. Inspired by eHPE (Chen et al. 2022a), the number c of events involved in the fusion is added as a new feature to the fused event point cloud. Unlike voxelization methods (Zhu et al. 2019; Rebecq et al. 2019) that use the time coordinates corresponding to voxel blocks, we retain the precise average time and the number of events involved in fusion, minimizing information loss to the extent possible. The number of temporal segments T during event fusion can be balanced between temporal precision and sampling loss based on different tasks.

Random sampling is used to select top M inputs from the fused event point cloud for later processing. Random sampling significantly reduces computational costs (Qi et al. 2017b) during sampling and does not require additional virtual points (Lin et al. 2024) for feature propagation. To reduce the interference caused by the differing numerical scales (pixels vs. seconds) in event point cloud computation, each input point $\mathbf{e}' = (h, w, t', p', c')$ is normalized as $\hat{\mathbf{e}} = (\frac{h}{H}, \frac{w}{W}, \frac{t'}{t_{max} - t_{min}}, p, c)$, where t_{max} and t_{min} are the maximum and minimum timestamps of the input point cloud. By normalizing the event feature scales, the aggregation stage helps to avoid feature imbalance due to scale differences, facilitating a more uniform extraction of event features across the spatiotemporal dimensions.

3.2 Spatio-Temporal Decoupling (STD)

Traditional point cloud networks often use Euclidean distance for neighbor indexing. However, the heavy spatiotemporal scale variation across different event data limits the applicability of Euclidean distance for feature aggregation. EP2T introduces a spatiotemporal weighted distance to mitigate the dilution of local feature diversity caused by spa-




Figure 3: **SFCA structure.** The input undergoes serialization using a space-filling curve (SFC), followed by an embedding layer. Subsequently, the data passes through L Encoder Layers (for S-SFCA and T-SFCA combinations, $L=2$; for ST-SFCA configurations, $L=4$), each of which integrates Grid Pool operations with dynamic patch rearrangement for adaptive feature grouping. The processed features then flow through 2 consecutive encoder blocks (E.B.s). The decoder layer finally generates the output.

tiotemporal heterogeneity. However, the need for task/data-specific weight tuning limits its generalization.

We argue that despite the imbalance in spatiotemporal scales, there still exists an inherent relationship between space and time. However, explicitly modeling this relationship often leads to coarse granularity and inefficiency. Therefore, we propose separating spatiotemporal information to obtain independent unimodal distance metrics, allowing the scale of neighboring points to adapt to the current modality and ensuring a balanced trade-off between sparsity and density. Subsequently, the local features extracted through the spatiotemporal decoupling feature extraction branch are further refined in the STA phase, where the network uses an attention module to fuse the spatiotemporal scaling relationship, thereby enabling diverse combinations of local features and enhancing their representational power. Specifically, given a set of fused, sampled, and normalized events $\hat{e} = (h, w, t', p', c')$ from EFS. As illustrated in Fig. 2, The sampled event point clouds are separately fed into temporal, spatial, and spatio-temporal SFCA modules. In S-SFCA, we solely use space as the distance metric to determine neighborhoods, conversely, T-SFCA uses time as the distance metric. ST-SFCA adopts the traditional Euclidean distance. We keep ST-SFCA with the decoupled branches since there is a strong inherent relationship between space and time. The decoupled spatiotemporal local features aim to optimize the collaborative expression of these features. Later ablation studies demonstrate the necessity and effectiveness of all three branches.

3.3 Space Filling Curve based Aggregation (SFCA)

Unlike conventional point clouds, event data requires large receptive fields to model multi-scale spatiotemporal correlations. While traditional methods derive semantics from precise boundary features, event streams encode semantic information through temporal-spatial evolution patterns. Previous methods are limited by the receptive fields: PointNet-based methods (Qi et al. 2017a,b) suffer from information discontinuity across sampling hierarchies and severe signal decay in deep layers; KNN-based methods (Zhao et al. 2021;

Wu et al. 2022) face computational bottlenecks (linear complexity with neighbor counts) when scaling to dense event streams.

Our work addresses this through a space-filling curve-driven module inspired by (Wu et al. 2024). This strategy enables efficient large-receptive-field learning by establishing continuous multi-scale correlations, directly aligning with the spatiotemporal nature of event data while maintaining computational feasibility.

Given a set of event points $P = \{(s'_i, p'_i, c'_i) \mid i = 1, 2, \dots, N\}$, where s' is the distance metric dimension(s) of the current SFCA branch, we first map it to a regular grid structure by applying $\lfloor s'_i/g \rfloor$ based on a given grid size g . This allows us to use the inverse mapping of a space-filling curve, to transform each point in s' into a one-dimensional encoding value as follows:

$$V_m = \phi_m \left(\left\lfloor \frac{s'_i}{g} \right\rfloor \right) \quad (2)$$

where $\phi_m(\cdot)$ represents different types of space-filling curves (e.g., the Z-Order curve (Morton 1966) or the Hilbert curve (Hilbert 1935)) and $V_m \in \mathbb{Z}$ denotes the 1-D index of point m on a space-filling curve. The event points P are then embedded into a feature representation $F^0 \in \mathbb{R}^{N \times C^0}$, which is then fed into the encoder layer.

To enhance the model’s generalization ability, in each encoder layer, we first randomly select a precomputed encoding value V_m from one of the space-filling curves and use it to sort the event points. The Hilbert curve is chosen as the encoding foundation due to its superior ability to preserve spatial continuity. By mapping 2D coordinates into a 1D sequence, it effectively maintains the spatial proximity of physically adjacent points. In contrast, the Z-order curve introduces spatial discontinuities due to its alternating coordinate encoding; however, this very property makes it well-suited for capturing non-contiguous event distributions.

In S-SFCA and T-SFCA, the expansion of the local receptive field does not rely on the jump-based connections of the Z-order curve. In fact, such discontinuous encoding could interfere with event feature extraction, which is why we exclusively adopt the Hilbert curve in these modules. However, in ST-SFCA, capturing temporal variations at the same pixel location proves challenging using the Hilbert curve alone, as its strong spatial continuity fails to account for long-range temporal dependencies. Here, the discontinuous connections of the Z-order curve become advantageous, allowing for the effective capture of local temporal variations across distant time steps. Therefore, we employ both the Hilbert and Z-order curves in our design.

Then, we partition F into $\lfloor N/p \rfloor$ non-overlapping patches $\hat{F}_j \in \mathbb{R}^{p \times C^0}$ ($j = 1, 2, \dots, \lfloor N/p \rfloor$), where $p = 512$ in our model, with each patch serving as a local neighborhood for feature extraction. This partitioning is achieved via encoding-based sorting with a linear time complexity of $O(N)$. In contrast, traditional KNN-based neighborhood search has an $O(NK)$ time/memory complexity where K is the number of neighboring points, e.g., 512 in our case. The immense computational burden makes it infeasible to signif-

icantly expand the receptive field, limiting previous methods to use only 16 points (Lin et al. 2024).

Each patch undergoes two self-attention encoder blocks. The first block transforms input patch features \hat{F}_j into intermediate representations \hat{F}_j^1 , while the second block refines them into $\hat{F}_j^2 \in \mathbb{R}^{P \times C^1}$, where C^1 is the output channel size of the first encoder layer.

Subsequently, the encoded features F^1 are fed into the next encoder layer. After selecting the V_m , we introduce a pooling operation before sorting. Specifically, we perform a right shift by y bits on the spatial encoding values of V_m , ensuring that originally adjacent points are mapped to the same value:

$$V'_m = V_m \gg y \quad (3)$$

When selecting y , unlike the fixed-value approach in (Wu et al. 2024), we account for the unique characteristics of event cameras. Events tend to be spatially dense but temporally sparse. This imbalance necessitates dynamic downsampling: we set $y = 5$ initially to reduce redundancy in densely clustered spatial points while preserving key structural features, then adjust it to $y = 3$ in later stages to balance the decreasing event density over time while maintaining a uniform distribution.

Based on this, we perform max pooling on points with the same V'_m value to achieve aggregation. This reduces the point count from N to $N^1 \approx N/2^y$, generating downsampled features $F_d^1 \in \mathbb{R}^{N^1 \times C^1}$. Then those downsampled points are processed similarly to the first layer, being fed into the encoder block for transformation.

After L layers hierarchical stacking, the model achieves exponential receptive field growth. The base coverage P defines the initial local perception, while 2^{yL} represents the cumulative expansion factor from L pooling layers. The total receptive field after L layers is:

$$R_L = P \cdot 2^{yL} \quad (4)$$

Theoretically, the final receptive field of SFCA would be 32 times bigger than K-NN based methods EP2T.

3.4 Spatio-temporal Separated Attention (STA) and Feature Tensorization (FT)

After completing feature aggregation, to further establish the correlation between time and space, we feed the features from the three SFCA branches into STA for feature computation based on attention mechanisms. The structure of STA is illustrated in Fig. 4, see Appendix A.1 for attention module architecture details.

The features F_s and F_t pass through the spatiotemporal cross-attention fusion module to capture the relationships between decoupled spatial-temporal dimensions. The enhanced features obtained from this process are then added residually to the input, forming F_s^1 and F_t^1 as the input for the next interaction stage. After four stages of spatiotemporal cross-attention fusion, F_s^4 and F_t^4 are globally fused with F_{st} in the Spatial-Temporal Interaction module. In this step, F_s^4 and F_t^4 serve as representative features in the spatial and temporal distance measurement domains to enhance the

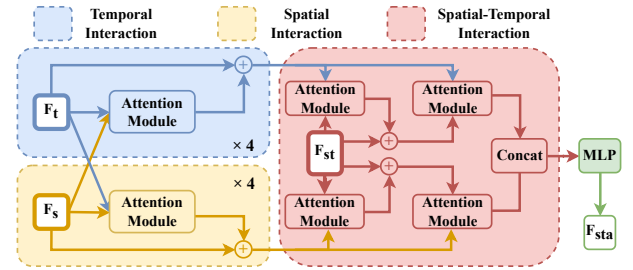


Figure 4: **STA structure.** STA integrates the temporal, spatial, and spatiotemporal features that were previously decoupled and encoded, completing the *decouple-enhance-fuse* paradigm. It enhances the temporal-only and spatial-only features from SFCA through four rounds of mutual attentions, and then refines these features by conducting multiple rounds of attentions with the spatiotemporal features. The refined features are fed into a multi-layer perceptron (MLP), yielding the output feature F_{sta} .

global feature F_{st} . Finally, the bidirectional enhancement results are concatenated, and the concatenated features are mapped through an MLP to obtain the output F_{sta} .

4 Experiments

To validate the applicability of OmniEvent, we conduct experiments on 3 representative and diverse tasks: classification (Sec. 4.1), optical flow estimation (Sec. 4.2) and event to point cloud registration (Sec. 4.3). Sec. 4.4 further performs detailed analysis, validating the effectiveness of individual components. All methods are implemented using PyTorch on an Intel Xeon Gold 6342 CPU and an NVIDIA A40 GPU.

4.1 Classification

Datasets. We conduct experiments on five event-based object classification datasets: N-Caltech101 (Orchard et al. 2015), N-MNIST (Orchard et al. 2015), CIFAR10-DVS (Li et al. 2017), N-CARS (Sironi et al. 2018), and N-ImageNet (Kim et al. 2021). The datasets N-Caltech101, N-MNIST, CIFAR10-DVS, and N-ImageNet are generated by recording moving images on an LCD screen using an event camera. Following the common practice, we keep 80% of the data for training and validation, and the remaining 20% for testing. We use the N-Caltech101 partition from (Gehrig et al. 2019; Cannici et al. 2020).

Implementation. Event representation learning focuses on feature extraction and representation transformation, so evaluating its performance requires selecting an appropriate task network. To ensure a fair comparison, we follow the settings of existing representation learning methods (Gehrig et al. 2019) and use a basic ResNet34 as the task network. Our SFCA (Fig. 3) consists of multiple layers of Sub-SFCA. Specifically, T-SFCA and S-SFCA each contain 2 layers, while ST-SFCA comprises 4 layers. Detailed architectural information for each layer can be found in the appendix. We follow EST (Gehrig et al. 2019) to use ADAM optimizer for all experiments with a learning rate of 10^{-4} . Due to mem-

Method	CIFAR10- DVS	N- Caltech101	N- CARS	N- MNIST	N- ImageNet
HATS	52.4	64.2	90.2	99.1	47.1
TORE	N/A	79.8	97.7	99.4	54.6
ECSNet	72.7	69.3	94.6	99.2	N/A
GraphSTL	54.0	65.7	91.4	99.0	N/A
VMV-GCN	69.0	77.8	93.2	99.5	N/A
MVF-Net	76.2	87.1	96.8	99.3	N/A
GET	<u>84.8</u>	N/A	96.7	<u>99.7</u>	N/A
General Event Representation Learning (ResNet34)					
EST	74.9	81.7	92.5	99.1	48.9
M-LSTM	73.1	84.1	95.8	98.9	32.2
EP2T	79.6	85.8	96.5	98.7	N/A
EventPillars	N/A	85.3	97.1	98.9	52.3
ERGO-12	74.0	77.2	92.3	99.2	61.4
OmniEvent(Ours)	85.2	90.2	97.9	99.9	62.1

Table 1: **Classification result.** We report top-1 classification accuracy, which is the best values from the original papers and our reproduced results. **Bold** and underlined values represent the best and second-best results. The upper part of the table shows SOTA task-specific methods, including HATS (Sironi et al. 2018), TORE (Baldwin et al. 2022), ECSNet, GraphSTL (Bi et al. 2020), VMV-GCN (Xie et al. 2022), MVF-Net (Deng, Chen, and Li 2021) and GET (Peng et al. 2023). The lower part shows general event representation learning methods with ResNet34, including EST, M-LSTM (Cannici et al. 2020) and EP2T, EventPillars (Fan et al. 2025) and ERGO-12 (Zubić et al. 2023).

ory constraints, we set the batch size to 10. We perform early stopping on a validation set in all experiments.

As shown in Tab. 1, OmniEvent outperforms both task-specific and general event-based methods. It is worth noting that OmniEvent, paired with a basic task network and devoid of any training hyperparameter tuning, achieves leading performance on *all* datasets. Also OmniEvent is the *only* method that is consistently better than others, the second best method varies across datasets, indicating limited generalization of existing frameworks.

4.2 Optical Flow Estimation

Datasets. We use the widely recognized Multi Vehicle Stereo Event Camera Dataset (MVSEC) (Zhu et al. 2018a) dataset for optical flow experiments. MVSEC contains indoor and outdoor scene sequences with significant differences and provides optical flow ground-truth.

Implementation. Following previous works (Gehrig et al. 2019; Cannici et al. 2020), we use the classic EV-FlowNet (Zhu et al. 2018b) (simple U-Net) as the task network, training on the *outdoor_day1*, and *outdoor_day2* sequences with optical flow ground truth as supervision. Following (Gehrig et al. 2019), we evaluate models by comparing the average end-point error ($AEE = \frac{1}{N} \sum_i |f_{est} - f_{gt}|_2$) and global outlier (where $AEE > 3$) ratio on the indoor flying sequences, which have different scenes and motion patterns from the training set to validate generalization. We also mask the hood area to avoid interference from reflections following E-cGAN (Mostafavi, Wang, and Yoon 2021). Additionally, due to the sparsity of events, we evaluate perfor-




Figure 5: **Qualitative comparisons.** The green box and the gray box represent the projection on the 2D view after the 3D point cloud is rigidly transformed to the camera coordinate system using the true pose and the predicted pose.

mance only at the pixel locations where events occur (sparse optical flow), as done in (Gehrig et al. 2019; Zhu et al. 2018b). Due to memory limitations, the batch size is set to 6 in optical flow estimation. Apart from the batch size, OmniEvent’s structure and other training parameters *remain the same* as in the classification tasks (Sec. 4.1).

As shown in Tab. 2, OmniEvent also surpasses all other representation learning methods. See Appendix B for visualization of optical flow estimation. It is worth noting that these specialized optical flow estimation networks incorporate additional information. For example, E-Distillation employs knowledge distillation from optical flow estimation networks based on traditional images. Despite using a simple U-Net as the task architecture *without hyperparameter tuning*, OmniEvent remains competitive against specialized optical flow estimation networks.

4.3 Event-to-Point Cloud Registration

Datasets. Following the previous SOTA (Lin et al. 2024), we evaluate the performance using the real-world datasets MVSEC (Zhu et al. 2018a) and VECtor (Gao et al. 2022). The train–test split follows (Lin et al. 2024).

Implementation. Following (Gehrig et al. 2019; Lin et al. 2024), we use the classic LCD (Pham et al. 2020) as the task network. Consistent with EP2T, we evaluate models using the translation error $TE = \|T_{GT} - T_{pred}\|_2$ and the rotation error $RE = \arccos\left(\frac{\text{tr}(R_{GT}^{-1}R_{pred}) - 1}{2}\right)$. Here, T and R denote the translation vector and rotation matrix. The notation $\text{tr}(\cdot)$ represents the trace of a matrix. The subscripts “GT” and “pred” indicate the ground-truth and predicted poses, respectively. The translation error TE is measured in meters, while the rotation error RE is expressed in degrees. We use the same hyperparameters as in the classification task except that the batch size is set to 4.

As shown in Tab. 3, OmniEvent outperforms previous SOTA by a large margin on both datasets. Compared with the best previous model EP2T, OmniEvent reduces the RE and TE by up to 68% and 60% respectively. Fig. 5 further visualizes the results of different methods. OmniEvent provides more accurate reprojections over SOTA methods.

Methods	indoor flying1		indoor flying2		indoor flying3	
	AEE(↓)	%Out.(↓)	AEE(↓)	%Out.(↓)	AEE(↓)	%Out.(↓)
EVFlowNet(EVF.)	1.03	2.20	1.72	15.10	1.53	11.90
E-Basics (Paredes-Vallés and De Croon 2021)	0.79	1.20	1.40	10.90	1.18	7.40
E-Distillation (Deng et al. 2021)	0.89	0.66	1.31	6.44	1.13	3.53
EST+EVF.	0.97	0.91	1.38	8.20	1.43	6.47
M-LSTM+EVF.	0.82	0.53	1.19	5.59	1.08	4.81
EP2T+EVF.	0.85	0.05	1.35	6.75	1.30	4.09
OmniEvent(Ours)+EVF.	0.70	0.03	1.12	<u>6.36</u>	1.15	3.33

Table 2: **Optical flow estimation result on MVSEC.** AEE represents the average end-point error and Out. represents the proportion of AEE errors exceeding 3 pixels. The upper part of the table presents task-specific methods, the bottom part compares representation learning methods combined with the EVFlowNet (Zhu et al. 2018b). **Bold** and underlined values represent the best and second-best results.

Input	Method	MVSEC-E2P		VECTor-E2P	
		RE(°)(↓)	TE(m)(↓)	RE(°)(↓)	TE(m)(↓)
Traditional Image	Grayscale Image	6.335	1.347	17.879	13.200
Event(Tensor-based)	E-Statistic	4.968	1.297	11.034	9.416
	TORÉ	4.855	1.350	9.521	<u>7.254</u>
	EP2T	3.606	0.821	8.672	7.403
	OmniEvent(Ours)	1.146	0.327	6.188	7.201

Table 3: **2D-3D registration result.** Since the 3 handcrafted methods (Zhu et al. 2018b,c; Mostafavi, Wang, and Yoon 2021) are mutually enhancing, we combine them to optimize performance, and call the combined method E-Statistic. Additionally, we include two other methods, TORÉ and EP2T, for a comprehensive evaluation. **Bold** and underlined values represent the best and second-best results.

No.	S-SFCA	ST-SFCA	T-SFCA	STA	Statistical Features	Accuracy (%)	Space (MB)
1	✓	✓	✓	✓	✓	85.2	3146
2		✓	✓	✓	✓	75.4	2600
3	✓		✓	✓	✓	79.2	2188
4	✓	✓		✓	✓	74.7	2654
5	✓	✓	✓		✓	76.7	2763
6	✓	✓	✓	✓		82.0	2596

Table 4: **Ablation.** Removing modules from OmniEvent hurts the classification accuracy on CIFAR10-DVS; the last column (Space) reports the GPU memory usage in MB.

4.4 Analysis

Ablation. To assess each module in OmniEvent, we perform ablations by removing individual components. For SFCA, we simply remove the corresponding branches. For STA, we replace it by feature concatenation from 3 SFCA branches. As shown in Tab. 4, the classification accuracy reduces as components are removed, proving the importance of each component, while the last column (Space) reports the GPU memory usage in MB.

Compared with the previous best model EP2T, one advantage of OmniEvent is the removal of manually tuned weights for different spatial-temporal branches, which simplifies training across tasks. As shown in Fig. 6 of Appendix A.3, EP2T requires different weight combinations to achieve optimal performance for different tasks/datasets, and even

though heavy weight tuning is conducted, EP2T still performs worse than OmniEvent for all datasets.

We provide further ablation results in Appendix C, including the importance of large neighborhood size in SFCA (Tab. 8), ablations on optical flow performance (Tab. 9), the type and combination of space-filling curves (Tab. 10), the number of layers in SFCA (Tab. 11), and the depth and attention mechanism in STA (Tab. 12 and 13).

Method	OmniEvent(Ours)	EP2T	ECSNet	EST
Time (ms)	36.6	42.8	67.8	6.3
Space (MB)	3146	4650	2530	2457
Accuracy (%)	85.2	<u>79.6</u>	72.7	74.9

Table 5: **Time and memory efficiency evaluated on CIFAR10-DVS.** **Bold** and underlined values represent the best and second-best results.

Memory and Time Efficiency. Besides the SOTA accuracy, OmniEvent also provides reasonable efficiency. As shown in Tab. 5, while EST achieves the fastest inference time, it sacrifices classification accuracy heavily, falling 10.3% behind our method. OmniEvent attains the highest accuracy (85.2%) with a competitive run time (36.6 ms), demonstrating a 14.4% speed improvement over EP2T and 45.9% faster than ECSNet. Meanwhile, OmniEvent requires only 3146MB memory, which is 1/3 lower than EP2T.

5 Conclusion

We propose OmniEvent, an innovative framework that can learn effective representations across various event camera tasks with the *same* architecture. OmniEvent proposes a novel decouple-enhance-fuse paradigm, which fully decouples spatial-temporal domains during local feature aggregation to handle spatial-temporal inhomogeneity, and learn global spatial-temporal relationships during the “fuse” stage with attention mechanisms. *Without* task-specific architecture design or hyperparameter tuning, OmniEvent advanced SOTA across 3 representative tasks and 10 datasets. While OmniEvent provides a unified architecture, our current experiments still train separate models for different tasks due to the limited dataset size. An interesting future direction is to scale up model training and produce a single foundation model for various tasks.

Acknowledgments

This work is supported in part by the National Natural Science Foundation of China (No. 62471415, 62401225), the Natural Science Foundation of Fujian Province of China (No. 2023J01004, 2025J0141, 2024J01115), the Natural Science Foundation of Xiamen of China (No. 3502Z202472018), and the Jimei University Scientific Research Start-up Funding Project (No.ZQ2024034).

References

- Almatrafi, M.; Baldwin, R.; Aizawa, K.; and Hirakawa, K. 2020. Distance surface for event-based optical flow. *IEEE TPAMI*, 42(7): 1547–1556.
- Andreopoulos, A.; Kashyap, H. J.; Nayak, T. K.; Amir, A.; and Flickner, M. D. 2018. A low power, high throughput, fully event-based stereo system. In *CVPR*, 7532–7542.
- Baldwin, R. W.; Liu, R.; Almatrafi, M.; Asari, V.; and Hirakawa, K. 2022. Time-ordered recent event (TORE) volumes for event cameras. *IEEE TPAMI*, 45(2): 2519–2532.
- Benosman, R.; Clercq, C.; Lagorce, X.; Ieng, S.-H.; and Bartolozzi, C. 2013. Event-based visual flow. *IEEE transactions on neural networks and learning systems*, 25(2): 407–417.
- Bi, Y.; Chadha, A.; Abbas, A.; Bourtsoulatze, E.; and Andreopoulos, Y. 2020. Graph-based spatio-temporal feature learning for neuromorphic vision sensing. *IEEE TIP*, 29: 9084–9098.
- Cannici, M.; Ciccone, M.; Romanoni, A.; and Matteucci, M. 2019. Asynchronous convolutional networks for object detection in neuromorphic cameras. In *CVPRW*, 0–0.
- Cannici, M.; Ciccone, M.; Romanoni, A.; and Matteucci, M. 2020. A differentiable recurrent surface for asynchronous event-based data. In *ECCV*, 136–152. Springer.
- Chen, J.; Shi, H.; Ye, Y.; Yang, K.; Sun, L.; and Wang, K. 2022a. Efficient human pose estimation via 3d event point cloud. In *2022 International Conference on 3D Vision (3DV)*, 1–10. IEEE.
- Chen, Z.; Wu, J.; Hou, J.; Li, L.; Dong, W.; and Shi, G. 2022b. ECSNet: Spatio-Temporal Feature Learning for Event Camera. *IEEE TCSVT*.
- Deng, Y.; Chen, H.; Chen, H.; and Li, Y. 2021. Learning from images: A distillation learning framework for event cameras. *IEEE TIP*, 30: 4919–4931.
- Deng, Y.; Chen, H.; and Li, Y. 2021. MVF-Net: A multi-view fusion network for event-based object classification. *IEEE TCSVT*, 32(12): 8275–8284.
- Fan, R.; Hao, W.; Guan, J.; Rui, L.; Gu, L.; Wu, T.; Zeng, F.; and Zhu, Z. 2025. EventPillars: Pillar-based Efficient Representations for Event Data. In *AAAI*, volume 39, 2861–2869.
- Gallego, G.; Delbrück, T.; Orchard, G.; Bartolozzi, C.; Taba, B.; Censi, A.; Leutenegger, S.; Davison, A. J.; Conrath, J.; Daniilidis, K.; et al. 2020. Event-based vision: A survey. *IEEE TPAMI*, 44(1): 154–180.
- Gao, L.; Liang, Y.; Yang, J.; Wu, S.; Wang, C.; Chen, J.; and Kneip, L. 2022. Vector: A versatile event-centric benchmark for multi-sensor slam. *IEEE Robotics and Automation Letters*, 7(3): 8217–8224.
- Gehrig, D.; Loquercio, A.; Derpanis, K. G.; and Scaramuzza, D. 2019. End-to-end learning of representations for asynchronous event-based data. In *ICCV*, 5633–5643.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hilbert, D. 1935. *Neubegründung der mathematik. erste mitteilung*. Springer.
- Huh, D.; and Sejnowski, T. J. 2018. Gradient descent for spiking neural networks. *NeurIPS*, 31.
- Kim, J.; Bae, J.; Park, G.; Zhang, D.; and Kim, Y. M. 2021. N-imagenet: Towards robust, fine-grained object recognition with event cameras. In *ICCV*, 2146–2156.
- Lagorce, X.; Orchard, G.; Galluppi, F.; Shi, B. E.; and Benosman, R. B. 2016. Hots: a hierarchy of event-based time-surfaces for pattern recognition. *IEEE TPAMI*, 39(7): 1346–1359.
- Li, H.; Liu, H.; Ji, X.; Li, G.; and Shi, L. 2017. Cifar10-dvs: an event-stream dataset for object classification. *Frontiers in neuroscience*, 11: 244131.
- Li, Y.; Bu, R.; Sun, M.; Wu, W.; Di, X.; and Chen, B. 2018. Pointcnn: Convolution on x-transformed points. *NeurIPS*, 31.
- Lin, X.; Qiu, C.; Shen, S.; Zang, Y.; Liu, W.; Bian, X.; Müller, M.; Wang, C.; et al. 2024. E2PNet: Event to Point Cloud Registration with Spatio-Temporal Representation Learning. *NeurIPS*, 36.
- Liu, Z.; Hu, H.; Cao, Y.; Zhang, Z.; and Tong, X. 2020. A closer look at local aggregation operators in point cloud analysis. In *ECCV*, 326–342. Springer.
- Manderscheid, J.; Sironi, A.; Bourdis, N.; Migliore, D.; and Lepetit, V. 2019. Speed invariant time surface for learning to detect corner points with event-based cameras. In *CVPR*, 10245–10254.
- Maqueda, A. I.; Loquercio, A.; Gallego, G.; García, N.; and Scaramuzza, D. 2018. Event-based vision meets deep learning on steering prediction for self-driving cars. In *CVPR*, 5419–5427.
- Mitrokhin, A.; Hua, Z.; Fermuller, C.; and Aloimonos, Y. 2020. Learning visual motion segmentation using event surfaces. In *CVPR*, 14414–14423.
- Morton, G. M. 1966. A computer oriented geodetic data base and a new technique in file sequencing.
- Mostafavi, M.; Wang, L.; and Yoon, K.-J. 2021. Learning to reconstruct hdr images from events, with applications to depth and flow prediction. *IJCV*, 129(4): 900–920.
- Nguyen, A.; Do, T.-T.; Caldwell, D. G.; and Tsagarakis, N. G. 2019. Real-time 6DOF pose relocalization for event cameras with stacked spatial LSTM networks. In *CVPRW*, 0–0.
- Orchard, G.; Jayawant, A.; Cohen, G. K.; and Thakor, N. 2015. Converting static image datasets to spiking neuromorphic datasets using saccades. *Frontiers in neuroscience*, 9: 159859.

- Paredes-Vallés, F.; and De Croon, G. C. 2021. Back to event basics: Self-supervised learning of image reconstruction for event cameras via photometric constancy. In *CVPR*, 3446–3455.
- Paredes-Vallés, F.; Scheper, K. Y.; and De Croon, G. C. 2019. Unsupervised learning of a hierarchical spiking neural network for optical flow estimation: From events to global motion perception. *IEEE TPAMI*, 42(8): 2051–2064.
- Park, P. K.; Cho, B. H.; Park, J. M.; Lee, K.; Kim, H. Y.; Kang, H. A.; Lee, H. G.; Woo, J.; Roh, Y.; Lee, W. J.; et al. 2016. Performance improvement of deep learning based gesture recognition using spatiotemporal demosaicing technique. In *ICIP*, 1624–1628. IEEE.
- Peng, Y.; Zhang, Y.; Xiong, Z.; Sun, X.; and Wu, F. 2023. GET: group event transformer for event-based vision. In *ICCV*, 6038–6048.
- Pham, Q.-H.; Uy, M. A.; Hua, B.-S.; Nguyen, D. T.; Roig, G.; and Yeung, S.-K. 2020. Lcd: Learned cross-domain descriptors for 2d-3d matching. In *AAAI*, volume 34, 11856–11864.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 30.
- Rebecq, H.; Ranftl, R.; Koltun, V.; and Scaramuzza, D. 2019. Events-to-video: Bringing modern computer vision to event cameras. In *CVPR*, 3857–3866.
- Rueckauer, B.; Lungu, I.-A.; Hu, Y.; Pfeiffer, M.; and Liu, S.-C. 2017. Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *Frontiers in neuroscience*, 11: 294078.
- Schaefer, S.; Gehrig, D.; and Scaramuzza, D. 2022. AEGNN: Asynchronous event-based graph neural networks. In *CVPR*, 12371–12381.
- Sekikawa, Y.; Hara, K.; and Saito, H. 2019. Eventnet: Asynchronous recursive event processing. In *CVPR*, 3887–3896.
- Sironi, A.; Brambilla, M.; Bourdis, N.; Lagorce, X.; and Benosman, R. 2018. HATS: Histograms of averaged time surfaces for robust event-based object classification. In *CVPR*, 1731–1740.
- Son, B.; Suh, Y.; Kim, S.; Jung, H.; Kim, J.-S.; Shin, C.; Park, K.; Lee, K.; Park, J.; Woo, J.; et al. 2017. 4.1 A 640×480 dynamic vision sensor with a 9μm pixel and 300Meps address-event representation. In *2017 IEEE International Solid-State Circuits Conference (ISSCC)*, 66–67. IEEE.
- Tavanaei, A.; Ghodrati, M.; Kheradpisheh, S. R.; Masquelier, T.; and Maida, A. 2019. Deep learning in spiking neural networks. *Neural networks*, 111: 47–63.
- Tulyakov, S.; Fleuret, F.; Kiefel, M.; Gehler, P.; and Hirsch, M. 2019. Learning an event sequence embedding for dense event-based deep stereo. In *ICCV*, 1527–1537.
- Wu, X.; Jiang, L.; Wang, P.-S.; Liu, Z.; Liu, X.; Qiao, Y.; Ouyang, W.; He, T.; and Zhao, H. 2024. Point Transformer V3: Simpler Faster Stronger. In *CVPR*, 4840–4851.
- Wu, X.; Lao, Y.; Jiang, L.; Liu, X.; and Zhao, H. 2022. Point transformer v2: Grouped vector attention and partition-based pooling. *NeurIPS*, 35: 33330–33342.
- Xie, B.; Deng, Y.; Shao, Z.; Liu, H.; and Li, Y. 2022. Vmv-gcn: Volumetric multi-view based graph cnn for event stream classification. *IEEE Robotics and Automation Letters*, 7(2): 1976–1983.
- Zhao, H.; Jiang, L.; Jia, J.; Torr, P. H.; and Koltun, V. 2021. Point transformer. In *ICCV*, 16259–16268.
- Zhou, Y.; Gallego, G.; Lu, X.; Liu, S.; and Shen, S. 2021. Event-based motion segmentation with spatio-temporal graph cuts. *IEEE Transactions on Neural Networks and Learning Systems*.
- Zhu, A. Z.; Thakur, D.; Özaslan, T.; Pfrommer, B.; Kumar, V.; and Daniilidis, K. 2018a. The multivehicle stereo event camera dataset: An event camera dataset for 3D perception. *IEEE Robotics and Automation Letters*, 3(3): 2032–2039.
- Zhu, A. Z.; Yuan, L.; Chaney, K.; and Daniilidis, K. 2018b. EV-FlowNet: Self-supervised optical flow estimation for event-based cameras. *arXiv preprint arXiv:1802.06898*.
- Zhu, A. Z.; Yuan, L.; Chaney, K.; and Daniilidis, K. 2018c. Unsupervised Event-based Learning of Optical Flow, Depth, and Egomotion. *arXiv:1812.08156*.
- Zhu, A. Z.; Yuan, L.; Chaney, K.; and Daniilidis, K. 2019. Unsupervised event-based learning of optical flow, depth, and egomotion. In *CVPR*, 989–997.
- Zubić, N.; Gehrig, D.; Gehrig, M.; and Scaramuzza, D. 2023. From chaos comes order: Ordering event representations for object recognition and detection. In *ICCV*, 12846–12856.