

Efficient Event-Based Semantic Segmentation via Exploiting Frame-Event Fusion: A Hybrid Neural Network Approach

Hebei Li¹, Yansong Peng¹, Jiahui Yuan¹, Peixi Wu¹, Jin Wang¹
Yueyi Zhang^{1,*}, Xiaoyan Sun^{1,2}

¹MoE Key Laboratory of Brain-inspired Intelligent Perception and Cognition, University of Science and Technology of China

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

{lihebei, pengyansong, yuanjiahui, wupeixi, jin01wang}@mail.ustc.edu.cn, {zhyueyi, sunxiaoyan}@ustc.edu.cn

Abstract

Event cameras have recently been introduced into image semantic segmentation, owing to their high temporal resolution and other advantageous properties. However, existing event-based semantic segmentation methods often fail to fully exploit the complementary information provided by frames and events, resulting in complex training strategies and increased computational costs. To address these challenges, we propose an efficient hybrid framework for image semantic segmentation, comprising a Spiking Neural Network branch for events and an Artificial Neural Network branch for frames. Specifically, we introduce three specialized modules to facilitate the interaction between these two branches: the Adaptive Temporal Weighting (ATW) Injector, the Event-Driven Sparse (EDS) Injector, and the Channel Selection Fusion (CSF) module. The ATW Injector dynamically integrates temporal features from event data into frame features, enhancing segmentation accuracy by leveraging critical dynamic temporal information. The EDS Injector effectively combines sparse event data with rich frame features, ensuring precise temporal and spatial information alignment. The CSF module selectively merges these features to optimize segmentation performance. Experimental results demonstrate that our framework not only achieves state-of-the-art accuracy across the DDD17-Seg, DSEC-Semantic, and M3ED-Semantic datasets but also significantly reduces energy consumption, achieving a 65% reduction on the DSEC-Semantic dataset.

Code —

<https://github.com/HebeiFast/Hybrid-Segmentation>

Introduction

Event cameras (Gallego et al. 2020), also known as neuromorphic cameras, are revolutionizing computer vision by offering a pathway to next-generation performance and efficiency. Unlike traditional frame-based cameras that capture images at fixed intervals, event cameras detect changes in pixel-level intensity asynchronously, focusing exclusively on the dynamic elements of visual scenes (Lichtsteiner, Posch, and Delbruck 2008). This unique capability provides high dynamic range and low power consumption (Gehrig and Scaramuzza 2022; Huang et al. 2023) and makes event

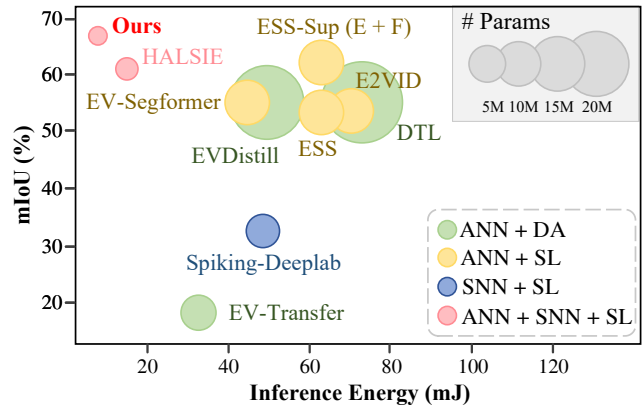


Figure 1: Comparison of various event-based semantic segmentation methods in terms of mIoU (%) and inference energy (mJ). DA and SL represent “domain adaptation” and “supervised learning”, respectively.

cameras particularly well-suited for applications such as autonomous driving, augmented reality, robotic vision, and other real-time processing tasks.

Among these applications, event-based semantic segmentation (ESS) stands out for its significance in visual perception (Alonso and Murillo 2019; Sun et al. 2022; Li et al. 2024a). However, existing ESS approaches often struggle to fully exploit the complementary information between frames and events. Most methods transfer information from the image domain to the event domain (Jing et al. 2024; Xia et al. 2023; Kong et al. 2024). However, these methods still require complex training strategies and incur significant computational costs due to the reliance on strong encoders.

To address these challenges, Spiking Neural Networks (SNNs) have recently emerged as a promising alternative for processing asynchronous data (Sengupta et al. 2019; Liao et al. 2023; Zhang et al. 2019; Yang et al. 2022). Unlike traditional artificial neural networks, SNNs operate with spiking neurons that transmit information via discrete spikes. This spiking mechanism offers several inherent advantages, including reduced energy usage, accelerated inference times, and a natural fit for handling event-driven tasks (Akopyan et al. 2015; Davies et al. 2018; Pei et al. 2019; Furber et al. 2014). Despite these advantages, recent studies that apply SNNs to segmentation tasks (Kim, Chough, and

*Corresponding Author

Panda 2022; Patel et al. 2021; Li et al. 2024b) indicate that there is still substantial room for performance enhancement. While some approaches have attempted to integrate ANN and SNN to leverage their complementary strengths (Biswas et al. 2024; Yue et al. 2023), these methods often neglect the critical interaction of intermediate layers between temporal information in SNNs and the spatial features extracted by ANNs, leaving significant potential untapped.

In this work, we propose a novel framework that seamlessly integrates the complementary characteristics of frame and event signals through two neural network branches. The ANN and SNN branches are responsible for feature extraction of the frames and events, respectively. In addition, three specialized modules are developed for the interaction between both branches: the Adaptive Temporal Weighting (ATW) Injector, the Event-Driven Sparse (EDS) Injector, and the Channel Selection Fusion (CSF) module. The ATW Injector is designed to dynamically incorporate temporal features from event data into frame features using adaptive temporal weighting, enhancing the frame features with crucial dynamic temporal information captured by event cameras, which significantly improves segmentation accuracy. The EDS Injector efficiently combines sparse event data with the rich spatial and color features of frames, ensuring precise alignment of temporal and spatial aspects critical for accurate scene understanding in dynamic environments. Additionally, the CSF module selectively integrates key semantic features from both branches, ensuring that only the most essential data contributes to the final segmentation. Experimental results demonstrate that our method achieves state-of-the-art results across the DDD17-Seg, DSEC-Semantic, and M3ED-Semantic datasets, while significantly reducing energy consumption (Figure 1). Notably, on the DSEC-Semantic dataset, our approach achieves a 65% reduction in energy consumption compared to the most efficient existing method.

In brief, our contributions are summarized as follows:

- We propose a novel framework that effectively integrates the complementary characteristics of frames and events, leveraging SNN and ANN in two branches to enhance event-based semantic segmentation.
- We introduce three components for branch interaction: the ATW Injector, which integrates dynamic temporal information from event data into the frame stream; the EDS Injector, which aligns sparse event features with frame features; and the CSF module, which selectively merges these features for optimal segmentation.
- Experimental results show that our method achieves state-of-the-art performance on multiple datasets, while also significantly reducing energy consumption.

Related Work

Event-based Semantic Segmentation

The goal of ESS is to classify pixels of frames into semantic categories with the help of events, thereby enhancing scene understanding. Ev-SegNet (Alonso and Murillo 2019) pioneered a new representation for event data and introduced the DDD17-Seg dataset specifically for ESS. To

reduce latency, HMNet (Hamaguchi et al. 2023) introduced a hierarchical memory-based event encoding approach. To improve accuracy while minimizing the need for extensive event annotations, EvDistil (Wang et al. 2021) and DTL (Wang, Chae, and Yoon 2021) leveraged knowledge distillation from labeled image data to train networks on unlabeled events. Ev-Transfer (Wang et al. 2021) and ESS (Sun et al. 2022) developed unsupervised domain adaptation techniques that adapt existing image datasets to the event domain. Addressing scalability and annotation challenges, OpenESS (Kong et al. 2024) proposed an annotation-free approach by distilling pre-trained vision-language models. More recently, HALSIE (Biswas et al. 2024) introduced an end-to-end learning framework for ESS that integrates cross-domain learning. In this work, we focus on processing cross-domain data while keeping energy efficiency in ESS.

Spiking Neural Network

SNNs, often regarded as the third generation of artificial neural networks, simulate biological neurons using binary spiking signals (Gerstner and Kistler 2002). Two main approaches for developing deep SNNs are ANN-to-SNN conversion and direct training. ANN-to-SNN conversion (Diehl et al. 2015; Rueckauer et al. 2017) translates trained ANNs into SNNs but is limited by the need for large time steps and poor handling of neuromorphic data. Direct training, which better suits temporal data, uses surrogate functions to overcome the non-differentiability of spiking neurons (Wu et al. 2018). However, these methods primarily target classification tasks and still underperform compared to ANNs, especially in complex tasks such as semantic segmentation.

To address the limitations of current models, several studies (Liu et al. 2023; Biswas et al. 2024) have investigated the combination of ANNs and SNNs. A notable approach employs an SNN for the encoder and an ANN for the decoder (Kugele et al. 2021; Zhang et al. 2021). This configuration optimizes performance by leveraging the precise temporal dynamics of SNNs while utilizing the robust representational capabilities of ANNs for reconstruction tasks. Additionally, multi-modal strategies (Lee, Kosta, and Roy 2022; Zhang et al. 2024) have been explored where RGB data is processed by ANNs and event data by SNNs, facilitating cross-modal fusion. This method significantly improves the system’s ability to interpret complex scenes with greater accuracy. Despite these advancements, most existing methods do not fully harness the potential of SNNs during feature extraction, often treating them merely as auxiliaries to ANNs without true integration. In this work, we aim to effectively integrate the continuous activations of ANNs with the discrete spikes of SNNs, thereby enhancing joint processing capabilities in complex perceptual tasks.

Preliminary

The original event stream comprises a sequence of 4-tuples $E = \{e_k\}_{k=1}^{N_e}$, where N_e denotes the total count of events. Each event e_k contains four attributes: (x_k, y_k, t_k, p_k) , which represent the spatial coordinates, timestamp, and the polarity of brightness change, respectively.

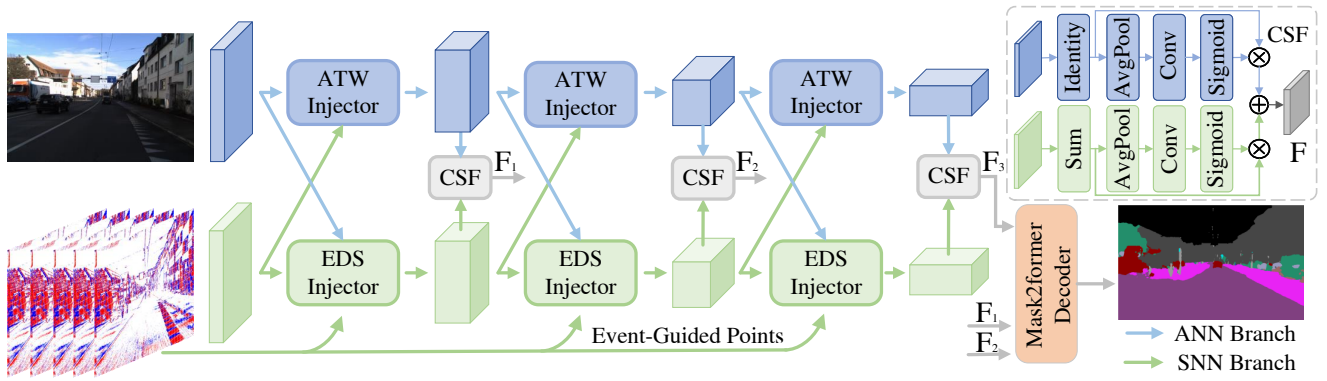


Figure 2: Overview of the proposed framework. The architecture integrates ANN and SNN branches for efficient semantic segmentation. The ATW Injector injects important event-based features within the ANN branch, while the EDS Injector enhances SNN features with color information from the ANN branch. The CSF module selectively fuses features from both branches. These multi-scale features from the CSF module are then fed into the Mask2Former Decoder.




Figure 3: Illustration of the proposed Adaptive Temporal Weighting (ATW) Injector.

In this work, we transform asynchronous event streams into corresponding voxel grids (Gehrig et al. 2019), represented as $V \in \mathbb{R}^{B \times H \times W}$, where B , H , and W denote the number of time bins, height, and width of the grids, respectively. Each time bin within the voxel grid is defined as

$$V(b) = \sum_i p_i \max(0, 1 - |b - \frac{t_i - t_{start}}{t_{end} - t_{start}}(B - 1)|), \quad (1)$$

where $b \in \{0, \dots, B-1\}$ represents the b -th time bin. t_{start} and t_{end} indicate the start and end timestamps of the event stream, respectively. To ensure uniformity in scale and mitigate the impact of outliers, we normalize the voxel grid V using the Z-score normalization.

Methodology

Main Architecture

The architecture of our method, illustrated in Figure 2, is built upon a novel frame-event fusion strategy within a hybrid neural network framework. This design integrates both ANN and SNN, leveraging the distinct advantages of conventional frame-based and event-based data. The architecture consists of two parallel branches: the ANN branch, which processes frame data using Adaptive Temporal Weighting (ATW) Injectors, and the SNN branch, which processes event data via Event-Driven Sparse (EDS) Injectors. Both branches converge in Channel Selection Fusion (CSF) modules, which extract the most relevant semantic features from each branch. These multi-scale features from the CSF modules are then fed into the Mask2Former Decoder, yielding the final segmentation output.

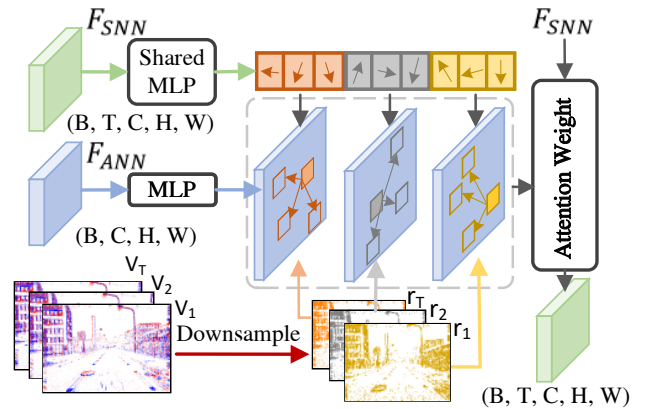


Figure 4: Illustration of the proposed Event-driven Sparse (EDS) Injector.

Adaptive Temporal Weighting Injector

The Adaptive Temporal Weighting (ATW) Injector, as illustrated in Figure 3, is designed to dynamically integrate SNN features into ANN features through cross attention. The primary objective is to leverage temporal information to optimize spatial feature representations. This integration not only improves segmentation accuracy in rapidly changing environments but also injects important event-based features into the ANN representation. The ATW initially receives an SNN feature from event data $F_{SNN} \in \{0, 1\}^{B \times T \times C \times H \times W}$. Considering the sparse nature of SNN features, we apply an average pooling operation to the spatial dimensions of F_{SNN} to effectively capture the overall sparsity across the $T \times C$, resulting in the squeezed feature $F_{SNN}^{Pool} \in \mathbb{R}^{B \times T \times C}$. Subsequently, we further utilize the Adaptor (Houlsby et al. 2019) to capture the global information along temporal and channel dimensions while maintaining low computational overhead. The adaptor contains a downsample and an upsample. The output of the adaptor is then averaged across the channel dimension C to produce a temporal weighting vector $W_{SNN} \in \mathbb{R}^{B \times T \times C}$

$$W_{SNN} = W_{up}^T \sigma(W_{down}^T F_{SNN}^{Pool}), \quad (2)$$

where W_{down} and W_{up} are the weights applied in the down-sample and upsample, respectively. σ is the ReLU activation function. Next, the dynamic weighted SNN features are obtained through the operation of normalizing the resulting temporal weights and multiplying them with the original SNN input tensor F_{SNN} , shown as

$$F_{SNN}^W = \sum_t^T \alpha_{t,c} \cdot F_{SNN}, \quad (3)$$

where $\alpha_{t,c}$ is the output of softmax function on W_{SNN} along channel and temporal dimension. Finally, this weighted tensor F_{SNN}^W is integrated with the ANN feature F_{ANN} via sparse attention layer (Zhu et al. 2020), described as

$$F_{ANN}^o = Attention(F_{SNN}^W, F_{ANN}), \quad (4)$$

where $F_{ANN}^o \in \mathbb{R}^{B \times C \times H \times W}$ is the optimized output feature map used in the subsequent stages of the ANN branch.

Event-driven Sparse Injector

The Event-driven Sparse (EDS) Injector, as illustrated in Figure 4, is designed to enhance the temporal feature extraction capabilities of SNN features by integrating them with the rich color information from ANN features derived from image data. By injecting detailed visual information into the sparse SNN features, the module enhances the SNN branch’s ability to efficiently and accurately extract temporal features, enabling a deeper understanding of complex dynamic scenes.

The EDS Injector module receives three distinct inputs: an SNN feature F_{SNN} , an ANN feature F_{ANN} and a voxel grid V . The SNN feature F_{SNN} , containing temporal changes in the scene, initially passes through a Shared MLP layer, which generates directional offsets for each time step T . These offsets adjust the spatial positioning of the temporal features. The resulting directional offsets are denoted as

$$\Delta r = ShareMLP(F_{SNN}), \quad (5)$$

where $\Delta r \in \mathbb{R}^{B \times T \times N_{offset} \times N_{points}}$ represents the different offset of the direction for sample points.

Previous methods (Zhu et al. 2020) use dense reference points during feature extraction, resulting in significant computational overhead for large feature maps. However, the sparsity of event signals does not require dense sampling. We leverage the sparse nature of event signals as guided reference points r to reduce computational costs. To obtain the sparse reference points r , we downsample the event voxel V to align the spatial resolution of F_{ANN} and acquire the non-zero values from the voxel grid as reference points. At the same time, the ANN features pass through an MLP layer, aligning them with the SNN feature space. The projected ANN features F_{ANN}^{Proj} are denoted as

$$F_{ANN}^{Proj} = MLP(F_{ANN}). \quad (6)$$

The projected ANN features F_{ANN}^{Proj} are then combined with the directional offsets Δr generated by the SNN branch. This combination uses deformable attention (Zhu et al. 2020) with event-guided reference points. The deformable

attention mechanism adjusts the sampling locations based on the directional offsets, allowing the network to integrate rich color information into SNN features

$$F_{SNN}^o = \sum_{k=1}^K A_{SNN} \cdot (F_{ANN}^{Proj} \cdot F_{SNN}[r_k + \Delta r_k]), \quad (7)$$

where A_{SNN} and F_{ANN}^{Proj} are the weighting coefficient and the projected ANN feature map, respectively. r_k , Δr_k and K are the reference positions from events, learned offset, and the number of sampling points, respectively. $F_{SNN}^o \in \mathbb{R}^{B \times T \times C \times H \times W}$ is the output feature of the EDS injector, delivered to subsequent stages of the SNN branch.

Channel Selection Fusion

The Channel Selection Fusion (CSF) module, as depicted in Figure 2, aims to effectively combine features from both the ANN and SNN branches to enhance the representation of critical semantic information. Inspired by (Hu, Shen, and Sun 2018), the CSF module emphasizes selecting the most informative channels from these two distinct feature sources.

The CSF module initially processes the F_{ANN}^o through an identity function to obtain x_{ANN} and the F_{SNN}^o through a sum operation to obtain x_{SNN} , respectively. After that, each branch utilizes the same operation flow for x_{ANN} and x_{SNN} , denoted as

$$s = x \cdot AvgPool(Conv(Sigmoid(x))) \quad (8)$$

where s and x represent the selected channel features and the input feature map from either the ANN or SNN branch. Following this, the weighted feature maps s_{ANN} and s_{SNN} are summed to produce the fused feature map $F \in \mathbb{R}^{B \times C \times H \times W}$. Through the channel selection and weighted fusion strategy, the CSF module highlights the feature channels by suppressing irrelevant or redundant features. Finally, the multi-scale outputs of the CSF module are sent to the mask2former decoder.

Experiments and Evaluation

Experimental Setup

For the implementation of our network, we incorporate frameworks from Detectron and SpikingJelly to establish a robust system for ESS. To facilitate smoother data handling and integration, all datasets in our study are standardized to a consistent format, ensuring consistency in data reading and loading. We train our model for 50000 iterations using the AdamW optimizer alongside the WarmupPolyLR scheduler. We also utilize AMP and gradient clipping techniques to optimize the network efficiently. For the SNN branch, we use the LIF model and the sigmoid function as surrogate gradients to approximate the gradients of the Heaviside function during backpropagation. The timestep of the SNN is 5. The tau and threshold parameters of LIF spiking neurons are 2.0 and 1.0, respectively. For event representations, the event window between consecutive frames is discretized into $B = 5$ temporal bins. These event representations and synchronized frames are processed by the ANN and SNN

Methods	Acc.[%]	mIoU[%]	Network	Para.(M)	#GFLOPs _{ANN}	#GFLOPs _{SNN}	E _{Total} (mJ)
EV-SegNet (Alonso and Murillo 2019)	89.76	54.81	ANN	29.09	73.62	-	338.65
EVDistill (Wang et al. 2021)	-	58.02	ANN	59.34	12.45	-	57.27
DTL (Wang, Chae, and Yoon 2021)	-	58.80	ANN	60.48	16.74	-	77.01
Spiking-Deeplab (Kim, Chough, and Panda 2022)	-	33.70	SNN	4.14	-	54.35	48.91
E2VID (Rebecq et al. 2019)	87.91	57.32	ANN	10.71	16.65	-	76.59
EV-Transfer (Messikommer et al. 2022)	47.37	14.91	ANN	7.37	7.88	-	36.25
ViD2E (Gehrig et al. 2020)	90.19	56.01	ANN	29.09	73.62	-	338.65
ESS (Sun et al. 2022)	88.43	53.09	ANN	12.91	14.22	-	65.41
ESS-Sup (E) (Sun et al. 2022)	91.08	61.37	ANN	12.91	14.22	-	65.41
ESS-Sup (E+F) (Sun et al. 2022)	90.37	60.43	ANN	12.91	14.22	-	65.41
EV-Segformer (Jia et al. 2023)	94.72	54.41	ANN	44.61	9.88	-	45.42
OpenESS (Kong et al. 2024)	91.05	63.00	ANN	-	-	-	-
HALSIE (Biswas et al. 2024)	92.50	60.66	Hybrid	1.82	3.84	0.267	17.89
Ours	95.07	67.31	Hybrid	1.79	1.95	0.110	9.08

Table 1: Quantitative segmentation results on the DDD17-Seg dataset.

DSEC-Semantic				
Methods	Acc.[%]	mIoU[%]	Para.(M)	E _{Total} (mJ)
EV-SegNet	88.61	51.76	29.09	1863.83
E2VID	87.55	49.66	10.71	416.99
EV-Transfer	60.50	23.20	7.37	197.48
ESS	84.16	45.38	12.91	356.32
ESS-Sup (E)	89.25	51.57	12.91	356.32
ESS-Sup (E+F)	89.37	53.29	12.91	356.32
OpenESS	90.21	57.21	-	-
HMNet-B1	88.94	51.42	4.07	1974.50
HMNet-L1	89.89	55.09	5.51	6686.38
HALSIE	89.01	52.43	1.82	94.41
Ours	94.27	66.57	1.79	33.45
M3ED-Semantic				
Methods	Acc.[%]	mIoU[%]	Para.(M)	E _{Total} (mJ)
E2VID	78.65	39.40	10.71	764.33
ESS-Sup (E)	75.97	34.82	12.91	471.56
ESS-Sup (E+F)	76.40	35.19	12.91	471.56
Ours	87.31	50.91	1.79	70.98

Table 2: Quantitative segmentation results on the DSEC-Semantic dataset and M3ED-Semantic dataset.

branches, respectively. In our experiments, we report accuracy and mean intersection over union (mIoU) on our segmentation maps to evaluate performance. We also estimate energy costs by calculating the average number of FLOPs for each dataset. Further details on the energy estimation process are found in the supplementary material. All experiments use 8 NVIDIA A800 GPUs.

Evaluation on DDD-17-Seg Dataset

Dataset and Training Details: DDD17-Seg (Alonso and Murillo 2019) is a well-established ESS benchmark dataset that consists of 40 sequences captured using a DAVIS346 camera along with synchronized grayscale frames, events and segmentation masks. The dataset includes a total of 15,950 training and 3,890 testing event sequences, each with a spatial resolution of 346×200 . For each time step, we con-

vert the 6400 events into an event voxel grid. To enhance the diversity of the training data, we apply data augmentation techniques such as resizing, cropping, and random flipping.

Results: Table 1 presents a comprehensive performance comparison of our method against several state-of-the-art approaches on the DDD17-Seg dataset. Our approach outperforms all competing methods, achieving an accuracy of 95.07% and a mIoU of 67.31%, which are the highest among the methods evaluated. Notably, our model utilizes an efficient network architecture, which significantly reduces the number of parameters to 1.79 million, and requires only 1.95 GFLOPs for the ANN component and 0.110 GFLOPs for the SNN component. This efficiency leads to a total energy consumption of just 9.08 mJ, making our method not only accurate but also computationally efficient. This demonstrates the effectiveness of our model in leveraging the strengths of both ANN and SNN architectures.

Rows 1-2 of Figure 5 provide qualitative comparisons of segmentation results on the DDD17-Seg dataset. Our approach consistently outperforms other methods, particularly in accurately segmenting challenging regions with complex driving scenes. The comparison between our method against E2VID and ESS-Sup (E) demonstrates that our method effectively captures finer details and exhibits improved segmentation accuracy, particularly in areas with rapid motion and fine structures. The visual results validate the efficacy of our framework in leveraging event-based information to enhance segmentation performance. Additionally, our method is able to capture semantic information about people that is not annotated in the ground truth.

Evaluation on DSEC-Semantic Dataset

Dataset and Training Details: The DSEC-Semantic dataset (Sun et al. 2022) extends the DSEC dataset by providing semantic labels for 11 sequences. This dataset includes 4017 training and 1395 testing event sequences, each with a spatial resolution of 640×440 . Different from DDD17-Seg, the number of selected events for each event voxel is 100000.

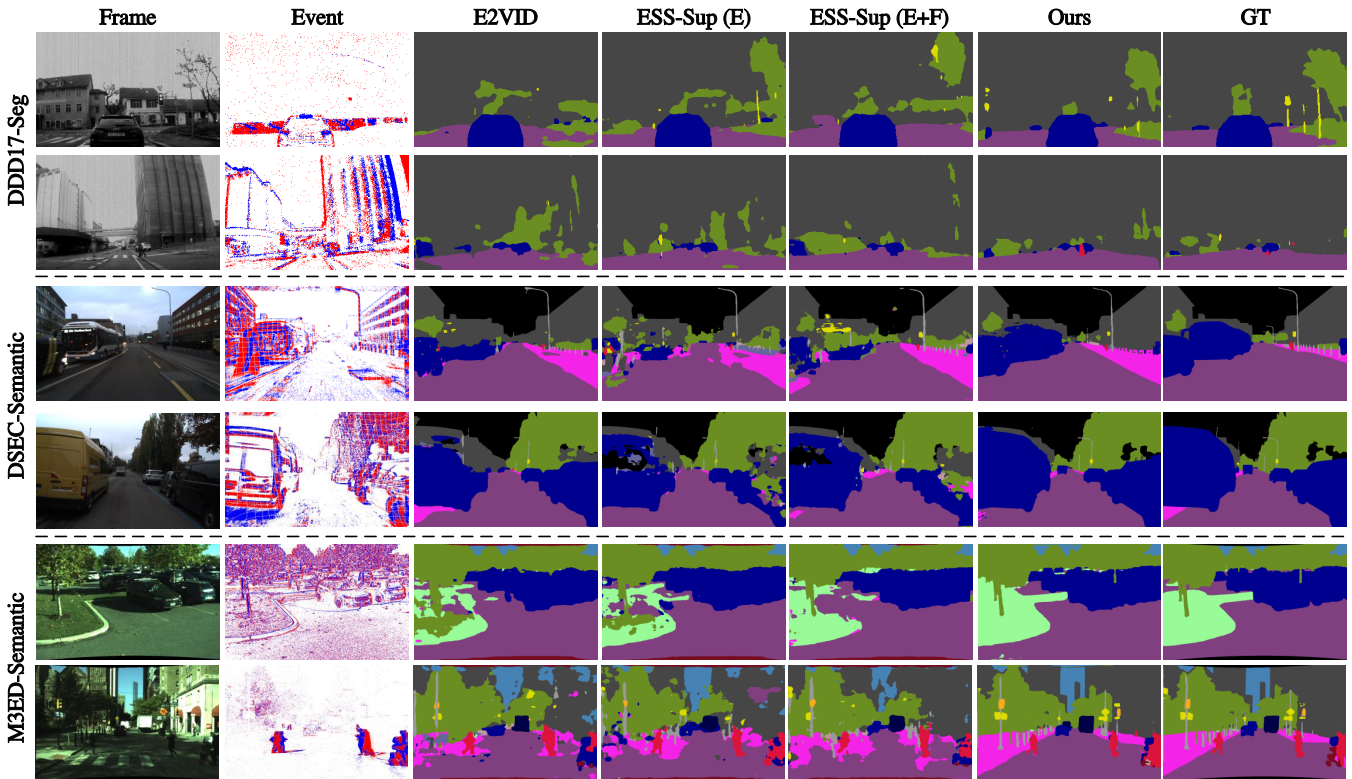


Figure 5: Qualitative segmentation results from the DDD17-Seg, DSEC-Semantic, and M3ED-Semantic datasets.

Results: Table 2 provides a detailed comparison of our method against various state-of-the-art approaches on the DSEC-Semantic dataset. Our method achieves the highest accuracy at 94.27% and a mIoU of 66.57%, outperforming all other methods. Despite its superior performance, our model is highly efficient, requiring only 1.79 million parameters, which is significantly lower than most other approaches. Additionally, our approach has a remarkably low energy consumption of 33.45 mJ, demonstrating the efficiency of our method in terms of energy resources. This efficiency is particularly notable when compared to methods such as EV-SegNet and HMNet-L1, which consume much higher energy levels. Overall, our approach balances high performance with low energy consumption, making it a robust solution for ESS tasks on the DSEC-Semantic dataset.

Rows 3-4 of Figure 5 provide qualitative comparisons of segmentation results on the DSEC-Semantic dataset. Our method demonstrates superior performance in capturing fine details and accurately segmenting complex urban scenes, particularly in areas with dynamic objects and challenging lighting conditions. Compared to other methods, our approach better aligns with the ground truth, effectively combining the temporal and spatial information to enhance semantic segmentation accuracy.

Evaluation on M3ED-Semantic Dataset

Dataset and Training Details: The M3ED-Semantic dataset, which is proposed in our paper, is a subset of the

M3ED dataset (Chaney et al. 2023) that provides large-scale high-quality, synchronized data from a variety of robotic platforms operating in diverse and challenging environments. We select 19 sequences, with 5,306 training sequences and 1,969 testing sequences for constructing the M3ED-Semantic dataset. Each sequence includes corresponding images, events, and segmentation masks annotated with 18 distinct labels. All data are captured at the resolution of 960×640 . The other setting is the same as the DSEC-Semantic dataset. The details description is provided in the supplementary material.

Table 2 presents a performance comparison of various methods on the M3ED-Semantic dataset, demonstrating the superiority of our approach in ESS. Our method achieves an accuracy of 87.31% and a mIoU of 50.91%, which are significantly higher than those of other competing methods. Specifically, E2VID records an accuracy of 78.65% and a mIoU of 39.40%, while the ESS-Sup (E) and ESS-Sup (E+F) methods report lower accuracies and mIoUs under similar conditions. Our method demonstrates remarkable energy efficiency with a total energy consumption of only 70.98 mJ, considerably less than the 764.33 mJ required by E2VID and 471.56 mJ by both ESS-Sup configurations. Overall, our method achieves high performance while maintaining low computational and energy costs.

Rows 5-6 of Figure 5 show the segmentation results from various methods on the M3ED-Semantic dataset. The results are illustrated for scenes captured from both drone

#	F-E	AWT	EDS	CSF	Acc.[%]	mIoU[%]
1	✓				92.82	61.06
2	✓	✓			93.23	63.59
3	✓		✓		93.05	63.71
4	✓			✓	93.01	62.98
5	✓	✓	✓		93.39	64.37
6	✓	✓		✓	93.58	64.18
7	✓		✓	✓	93.60	64.36
8	✓	✓	✓	✓	94.27	66.57

Table 3: Performance of models with different configurations on the DSEC-Semantic dataset.

Fusion	Acc.[%]	mIoU[%]	Para. (M)
EOLO (Cao et al. 2024)	93.39	64.39	4.50
REF (Zhou et al. 2023)	93.53	65.14	2.98
CF (Jiang et al. 2024)	93.49	64.19	6.32
Ours	94.27	66.57	1.79

Table 4: Quantitative results of different fusion methods on the DSEC-Semantic dataset.

and vehicle-mounted cameras. Our method exhibits superior segmentation performance, particularly in distinguishing between different classes in complex outdoor environments, such as roads, buildings, and vegetation. Compared to other methods, our approach demonstrates better alignment with the ground truth, accurately segmenting key elements in the scene even under challenging conditions like varying viewpoints and lighting. The integration of event data with frame-based features significantly enhances the model’s ability to capture fine details, leading to more precise segmentation in dynamic and diverse scenarios.

Ablation Study

The Effectiveness of AWT, EDS, and CSF Modules. Table 3 presents the impact of different module configurations on the performance of our model. The baseline model, consisting of frame and event branches, achieves an accuracy of 92.82% and a mIoU of 61.06%. When the AWT module is added to the baseline, the performance significantly improves, with accuracy rising to 93.23% and mIoU to 63.59%. This improvement underscores the importance of dynamically integrating temporal information into the RGB feature extraction process. Further enhancement is observed with the inclusion of the EDS module alongside AWT, resulting in an accuracy of 93.39% and mIoU of 64.37%. The EDS module effectively refines the event data processing by focusing on dynamically sampled features from RGB data, which contributes to the overall accuracy and segmentation quality. The addition of the CSF module, which refines feature representations on a channel-wise basis, achieves the highest accuracy of 94.27% and mIoU of 66.57%.

Evaluation of Fusion Strategies. Table 4 evaluates the effectiveness of different fusion strategies in our model. The EOLO method achieves an accuracy of 93.39% and a mIoU of 64.39%, with a parameter count of 4.50M. REF achieves

Timesteps	1	3	5	7
Accuracy [%]	93.69	93.91	94.27	93.88
mIoU [%]	64.90	65.45	66.57	65.14

Table 5: Performance of models at different timesteps on the DSEC-Semantic dataset.

slightly higher accuracy and mIoU scores at 93.53% and 65.14% with a reduced parameter count of 2.98M, indicating a more efficient use of parameters. The CF method, while comparable in accuracy to EOLO, slightly underperforms in mIoU and requires a significantly higher number of parameters. Our proposed fusion method outperforms all others while using only 1.79M parameters, the lowest among all the methods compared. This result demonstrates the superior efficiency and effectiveness of our fusion strategy, which improves performance and significantly reduces the complexity in terms of parameters.

Evaluation of Timesteps. Table 5 illustrates the impact of varying the number of timesteps on our model’s performance. Initially, using a single timestep results in an accuracy of 93.69% and a mIoU of 64.90%. As the number of timesteps increases to 3, there are slight improvements, with accuracy rising to 93.91% and mIoU to 65.45%. The optimal performance is achieved with 5 timesteps. This indicates that the model benefits from a balanced integration of temporal information in this configuration. However, when the timesteps are further increased to 7, both accuracy and mIoU show a slight decline, suggesting that the model may struggle to effectively utilize excessive temporal data, which could introduce noise or increase computational complexity, ultimately limiting performance.

Conclusion

In this paper, we introduced a novel hybrid neural network framework that integrates ANN and SNN to enhance the performance and efficiency of event-based semantic segmentation tasks. By leveraging the complementary strengths of both ANNs and SNNs, our framework effectively processes frame-based and event-based data, addressing the challenges of existing event-based semantic segmentation methods. The key components of our framework include the ATW Injector, which optimizes temporal feature extraction; the EDS Injector, which aligns and integrates sparse event data with rich spatial and color features; and the CSF module, which selectively merges features from both branches. Experimental results across multiple datasets demonstrate that our method achieves state-of-the-art segmentation accuracy and significantly reduces energy consumption. This work paves the way for more advanced hybrid models that can further optimize the synergy between frame-based and event-based processing for various computer vision tasks.

Acknowledgments

This work was in part supported by the National Natural Science Foundation of China under grants 62472399 and

References

- Akopyan, F.; Sawada, J.; Cassidy, A.; Alvarez-Icaza, R.; Arthur, J.; Merolla, P.; Imam, N.; Nakamura, Y.; Datta, P.; Nam, G.-J.; et al. 2015. TrueNorth: Design and Tool Flow of a 65 mW 1 Million Neuron Programmable Neurosynaptic Chip. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34(10): 1537–1557.
- Alonso, I.; and Murillo, A. C. 2019. EV-SegNet: Semantic segmentation for event-based cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*.
- Biswas, S. D.; Kosta, A.; Liyanagedera, C.; Apolinario, M.; and Roy, K. 2024. Halsie: Hybrid approach to learning segmentation by simultaneously exploiting image and event modalities. In *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 5952–5962. IEEE.
- Cao, J.; Zheng, X.; Lyu, Y.; Wang, J.; Xu, R.; and Wang, L. 2024. Chasing day and night: Towards robust and efficient all-day object detection guided by an event camera. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 9026–9032. IEEE.
- Chaney, K.; Cladera, F.; Wang, Z.; Bisulco, A.; Hsieh, M. A.; Korpela, C.; Kumar, V.; Taylor, C. J.; and Daniilidis, K. 2023. M3ed: Multi-robot, multi-sensor, multi-environment event dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4016–4023.
- Davies, M.; Srinivasa, N.; Lin, T.-H.; Chinya, G.; Cao, Y.; Choday, S. H.; Dimou, G.; Joshi, P.; Imam, N.; Jain, S.; et al. 2018. Loihi: A Neuromorphic Manycore Processor with On-Chip Learning. *IEEE Micro*, 38(1): 82–99.
- Diehl, P. U.; Neil, D.; Binas, J.; Cook, M.; Liu, S.-C.; and Pfeiffer, M. 2015. Fast-Classifying, High-Accuracy Spiking Deep Networks Through Weight and Threshold Balancing. In *2015 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Furber, S. B.; Galluppi, F.; Temple, S.; and Plana, L. A. 2014. The SpiNNaker Project. *Proceedings of the IEEE*, 102(5): 652–665.
- Gallego, G.; Delbrück, T.; Orchard, G.; Bartolozzi, C.; Taba, B.; Censi, A.; Leutenegger, S.; Davison, A. J.; Conradt, J.; Daniilidis, K.; et al. 2020. Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 44(1): 154–180.
- Gehrig, D.; Gehrig, M.; Hidalgo-Carrió, J.; and Scaramuzza, D. 2020. Video to events: Recycling video datasets for event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3586–3595.
- Gehrig, D.; Loquercio, A.; Derpanis, K. G.; and Scaramuzza, D. 2019. End-to-end learning of representations for asynchronous event-based data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 5633–5643.
- Gehrig, D.; and Scaramuzza, D. 2022. Are high-resolution event cameras really needed? *arXiv preprint arXiv:2203.14672*.
- Gerstner, W.; and Kistler, W. M. 2002. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press.
- Hamaguchi, R.; Furukawa, Y.; Onishi, M.; and Sakurada, K. 2023. Hierarchical neural memory network for low latency event processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 22867–22876.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *International Conference on Machine Learning (ICML)*, 2790–2799. PMLR.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 7132–7141.
- Huang, K.; Zhang, S.; Zhang, J.; and Tao, D. 2023. Event-based simultaneous localization and mapping: A comprehensive survey. *arXiv preprint arXiv:2304.09793*.
- Jia, Z.; You, K.; He, W.; Tian, Y.; Feng, Y.; Wang, Y.; Jia, X.; Lou, Y.; Zhang, J.; Li, G.; et al. 2023. Event-based semantic segmentation with posterior attention. *IEEE Transactions on Image Processing (TIP)*, 32: 1829–1842.
- Jiang, J.; Zhou, X.; Wang, B.; Deng, X.; Xu, C.; and Shi, B. 2024. Complementing event streams and rgb frames for hand mesh reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 24944–24954.
- Jing, L.; Ding, Y.; Gao, Y.; Wang, Z.; Yan, X.; Wang, D.; Schaefer, G.; Fang, H.; Zhao, B.; and Li, X. 2024. HPL-ESS: Hybrid Pseudo-Labeling for Unsupervised Event-based Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 23128–23137.
- Kim, Y.; Chough, J.; and Panda, P. 2022. Beyond classification: Directly training spiking neural networks for semantic segmentation. *Neuromorphic Computing and Engineering*, 2(4): 044015.
- Kong, L.; Liu, Y.; Ng, L. X.; Cottureau, B. R.; and Ooi, W. T. 2024. OpenESS: Event-based Semantic Scene Understanding with Open Vocabularies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15686–15698.
- Kugele, A.; Pfeil, T.; Pfeiffer, M.; and Chicca, E. 2021. Hybrid SNN-ANN: Energy-efficient classification and object detection for event-based vision. In *DAGM German Conference on Pattern Recognition*, 297–312. Springer.
- Lee, C.; Kosta, A. K.; and Roy, K. 2022. Fusion-FlowNet: Energy-efficient optical flow estimation using sensor fusion and deep fused spiking-analog network architectures. In *2022 International Conference on Robotics and Automation (ICRA)*, 6504–6510. IEEE.

- Li, H.; Wang, J.; Yuan, J.; Li, Y.; Weng, W.; Peng, Y.; Zhang, Y.; Xiong, Z.; and Sun, X. 2024a. Event-assisted Low-Light Video Object Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3250–3259.
- Li, H.; Zhang, Y.; Xiong, Z.; and Sun, X. 2024b. Deep multi-threshold spiking-UNet for image processing. *Neurocomputing*, 586: 127653.
- Liao, X.; Wu, Y.; Wang, Z.; Wang, D.; and Zhang, H. 2023. A Convolutional Spiking Neural Network with Adaptive Coding for Motor Imagery Classification. *Neurocomputing*, 126470.
- Lichtsteiner, P.; Posch, C.; and Delbruck, T. 2008. A 128×128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-state Circuits*, 43(2): 566–576.
- Liu, Z.; Wu, J.; Shi, G.; Yang, W.; Dong, W.; and Zhao, Q. 2023. Motion-oriented hybrid spiking neural networks for event-based motion deblurring. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*.
- Messikommer, N.; Gehrig, D.; Gehrig, M.; and Scaramuzza, D. 2022. Bridging the gap between events and frames through unsupervised domain adaptation. *IEEE Robotics and Automation Letters (RAL)*, 7(2): 3515–3522.
- Patel, K.; Hunsberger, E.; Batir, S.; and Eliasmith, C. 2021. A spiking neural network for image segmentation. *arXiv preprint arXiv:2106.08921*.
- Pei, J.; Deng, L.; Song, S.; Zhao, M.; Zhang, Y.; Wu, S.; Wang, G.; Zou, Z.; Wu, Z.; He, W.; et al. 2019. Towards artificial general intelligence with hybrid Tianjic chip architecture. *Nature*, 572(7767): 106–111.
- Rebecq, H.; Ranftl, R.; Koltun, V.; and Scaramuzza, D. 2019. High speed and high dynamic range video with an event camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 43(6): 1964–1980.
- 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: 682.
- Sengupta, A.; Ye, Y.; Wang, R.; Liu, C.; and Roy, K. 2019. Going Deeper in Spiking Neural Networks: VGG and Residual Architectures. *Frontiers in Neuroscience*, 13: 95.
- Sun, Z.; Messikommer, N.; Gehrig, D.; and Scaramuzza, D. 2022. Ess: Learning event-based semantic segmentation from still images. In *European Conference on Computer Vision (ECCV)*, 341–357. Springer.
- Wang, L.; Chae, Y.; and Yoon, K.-J. 2021. Dual transfer learning for event-based end-task prediction via plug-gable event to image translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2135–2145.
- Wang, L.; Chae, Y.; Yoon, S.-H.; Kim, T.-K.; and Yoon, K.-J. 2021. Evdistill: Asynchronous events to end-task learning via bidirectional reconstruction-guided cross-modal knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 608–619.
- Wu, Y.; Deng, L.; Li, G.; Zhu, J.; and Shi, L. 2018. Spatio-Temporal Backpropagation for Training High-Performance Spiking Neural Networks. *Frontiers in Neuroscience*, 12: 331.
- Xia, R.; Zhao, C.; Zheng, M.; Wu, Z.; Sun, Q.; and Tang, Y. 2023. Cmda: Cross-modality domain adaptation for nighttime semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 21572–21581.
- Yang, Q.; Wu, J.; Zhang, M.; Chua, Y.; Wang, X.; and Li, H. 2022. Training Spiking Neural Networks with Local Tandem Learning. *36rd Conference on Neural Information Processing Systems (NIPS)*.
- Yue, Y.; Baltés, M.; Abujahar, N.; Sun, T.; Smith, C. D.; Bihl, T.; and Liu, J. 2023. Hybrid Spiking Neural Networks Fine-Tuning for Hippocampus Segmentation. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, 1–5. IEEE.
- Zhang, L.; Zhou, S.; Zhi, T.; Du, Z.; and Chen, Y. 2019. TDSNN: From Deep Neural Networks to Deep Spike Neural Networks with Temporal-Coding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1319–1326.
- Zhang, R.; Leng, L.; Che, K.; Zhang, H.; Cheng, J.; Guo, Q.; Liao, J.; and Cheng, R. 2024. Accurate and efficient event-based semantic segmentation using adaptive spiking encoder-decoder network. *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*.
- Zhang, X.; Liao, W.; Yu, L.; Yang, W.; and Xia, G.-S. 2021. Event-based synthetic aperture imaging with a hybrid network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14235–14244.
- Zhou, Z.; Wu, Z.; Boutteau, R.; Yang, F.; Demonceaux, C.; and Ginjac, D. 2023. Rgb-event fusion for moving object detection in autonomous driving. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 7808–7815. IEEE.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable DETR: Deformable Transformers for End-to-End Object Detection. In *International Conference on Learning Representations (ICLR)*.