



**HAL**  
open science

## Event-Free Moving Object Segmentation from Moving Ego Vehicle

Zhuyun Zhou, Zongwei Wu, Danda Pani Paudel, Rémi Boutteau, Fan Yang,  
Luc Van Gool, Radu Timofte, Dominique Ginjac

► **To cite this version:**

Zhuyun Zhou, Zongwei Wu, Danda Pani Paudel, Rémi Boutteau, Fan Yang, et al.. Event-Free Moving Object Segmentation from Moving Ego Vehicle. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2024), Oct 2024, Abu Dhabi, United Arab Emirates. hal-04763949

**HAL Id: hal-04763949**

<https://hal.science/hal-04763949v1>

Submitted on 28 Jan 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Event-Free Moving Object Segmentation from Moving Ego Vehicle

Zhuyun Zhou<sup>1,2</sup>, Zongwei Wu<sup>2\*</sup>, Danda Pani Paudel<sup>3</sup>, Rémi Boutheau<sup>5</sup>, Fan Yang<sup>1</sup>,  
Luc Van Gool<sup>3,4</sup>, Radu Timofte<sup>2</sup>, and Dominique Ginjac<sup>1</sup>

**Abstract**—Moving object segmentation (MOS) in dynamic scenes is an important, challenging, but under-explored research topic for autonomous driving, especially for sequences obtained from moving ego vehicles. Most segmentation methods leverage motion cues obtained from optical flow maps. However, since these methods are often based on optical flows that are pre-computed from successive RGB frames, this neglects the temporal consideration of events occurring within the inter-frame, consequently constraining its ability to discern objects exhibiting relative staticity but genuinely in motion. To address these limitations, we propose to exploit event cameras for better video understanding, which provide rich motion cues without relying on optical flow. To foster research in this area, we first introduce a novel large-scale dataset called DSEC-MOS for moving object segmentation from moving ego vehicles, which is the first of its kind. For benchmarking, we select various mainstream methods and rigorously evaluate them on our dataset. Subsequently, we devise EmoFormer, a novel network able to exploit the event data. For this purpose, we fuse the event temporal prior with spatial semantic maps to distinguish genuinely moving objects from the static background, adding another level of dense supervision around our object of interest. Our proposed network relies only on event data for training but does not require event input during inference, making it directly comparable to frame-only methods in terms of efficiency and more widely usable in many application cases. The exhaustive comparison highlights a significant performance improvement of our method over all other methods. The source code and dataset are publicly available at: <https://github.com/ZZY-Zhou/DSEC-MOS>.

## I. INTRODUCTION

Achieving precise moving object segmentations (MOS) in urban scenes plays a vital role in many vision tasks like image/video synthesis [1], [2]. This segmentation is always achieved through cross-frame affinity matrices [3], [4]. However, in more challenging scenarios, such as videos recorded

This research is financed in part by the French National Research Agency through ANR CERBERE (ANR-21-CE22-0006), by the Ministry of Education and Science of Bulgaria (support for INSAIT, part of the Bulgarian National Roadmap for Research Infrastructure), and by the Alexander von Humboldt Foundation.

<sup>1</sup>Zhuyun Zhou, Fan Yang, and Dominique Ginjac are with University of Burgundy, Dijon, France. {Zhuyun.Zhou@etu., fanyang@, dginjac@}u-bourgogne.fr

<sup>2</sup>Zhuyun Zhou, Zongwei Wu, and Radu Timofte are with University of Wurzburg, Wurzburg, Germany. firstname.lastname@uni-wuerzburg.de

<sup>3</sup>Danda Pani Paudel and Luc Van Gool are with INSAIT, Sofia, Bulgaria. firstname.lastname@insait.ai

<sup>4</sup>Luc Van Gool is also affiliated with CVL, ETH Zurich, Zurich, Switzerland. vangool@vision.ee.ethz.ch

<sup>5</sup>Rémi Boutheau is with Université Rouen Normandie, INSA Rouen Normandie, Université Le Havre Normandie, Normandie Université, LITIS UR 4108, Rouen, France. remi.boutheau@univ-rouen.fr

\* Corresponding author

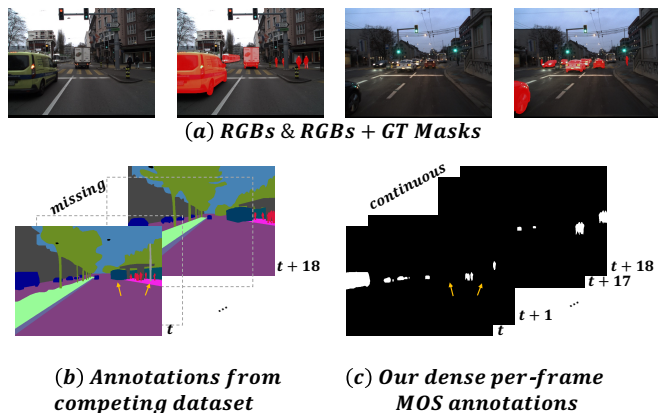


Fig. 1: DSEC-MOS Examples and Visual Comparison. The top part (a) shows calibrated-to-Event RGB frames, and our DSEC-MOS Ground Truth Segmentation Masks visualized on calibrated RGB frames. The bottom part (b, c) shows that our dataset provides per-frame annotation and distinguishes the motion attributes, which are not available in the previous dataset [11]. Best zoomed in.

from moving ego vehicles, conventional methods often falter, primarily due to the significant impact of unavoidable ego-motion [5], [6]. Consequently, there exists a critical need for a high-performing MOS model from moving ego vehicles.

Recently, event cameras [7], [8] have disrupted traditional computer vision paradigms. Different from frame-based cameras, event cameras operate asynchronously [7], enabling them to capture rapid motion and react to dynamic scenes with remarkable precision. In addition, their sensitivity to pixel intensity changes and innate ability to adapt to varying illumination conditions make them a robust choice for challenging scenarios [9], [10].

In this work, we aim to fully exploit the potential of event cameras for MOS from moving ego vehicles. Our problem extends the conventional (automatic) video object segmentation (VOS) [12] in a broader, challenging, and more daily-usage oriented setting with three crucial differences: (a) VOS typically aims to segment salient or foreground objects [13], [14], often a single object, while our goal is to segment all moving objects at once; (b) most automatic VOS methods preprocess the optical flow as input [15], [16] while we leverage ground truth event data; (c) in VOS, the camera motion is usually negligible compared to the object motion, while in our scenarios both objects and the cameras (embedded in the ego vehicle) are in unknown motion. These differences make our task an exceptional challenge, which, to our knowledge, is the first time being addressed

comprehensively.

To promote the research in the domain, we introduce a substantial dataset of 13,314 frames, derived from the DSEC dataset [17], [18]. This new dataset, called DSEC-MOS, covers a wide range of challenging driving scenarios, including different lighting conditions and complex traffic scenarios, as shown in Figure 1. It includes 8 categories of moving objects such as vehicles, pedestrians, cyclists, etc., each with different motion patterns, appearances, and scales, making the dataset highly suitable for both training and evaluation of MOS algorithms. Furthermore, DSEC-MOS provides dense pixel-level annotations for moving objects.

Leveraging our proposed dataset, we performed a systematic benchmark involving mainstream object segmentation methods in video tasks. Intriguingly, we observe that the leading methods [14], [15], [19], [20] do not necessarily maintain their competitiveness on our challenging dataset. Nevertheless, with the same backbone, RGB-flow methods [14], [19] consistently outperform RGB-only approaches [15], [21]. Additionally, [14], [20] with deeper architectures such as transformers [22] outperform CNN-based counterparts [23], [24] even when the latter uses extra flow cues. Consequently, we introduce a straightforward yet highly effective model, termed EmoFormer, that incorporates event clues as additional dense supervision during training while remaining event-free during inference. Detailed comparative analyses show that our method significantly outperforms all existing SOTA approaches.

In summary, our contributions are three-fold:

- We introduce a new MOS task with moving ego vehicles and present a generously sized and densely annotated dataset called DSEC-MOS.
- To establish a benchmark, we rigorously select mainstream segmentation methods in video tasks with pure image input or with image-flow inputs.
- We introduce EmoFormer, which fully benefits from event clues during training while being event-free during inference. Our model outperforms all SOTA counterparts by a large margin and can serve as a strong baseline to facilitate further research in this domain.

## II. RELATED WORK

**Automatic Video Object Segmentation (AVOS):** AVOS represents a specialized task focused on segmenting objects within video sequences without prior knowledge about the target objects [12]. Within the realm of AVOS, the primary cues for segmentation stem from the motion exhibited by objects. Traditional approaches involve computing similarity matrices [25]–[27] across different video frames, implicitly harnessing temporal context to guide object identification. Another category of AVOS methods [28]–[32] explicitly uses the optical flow as a guiding cue. Considering that target objects in AVOS are typically salient and predominantly occupy the foreground, and camera motion is generally negligible in comparison to object motion [33]–[35], extracting the target object from the flow map is relatively straightforward compared to a real-world setting [36].

**Event Vision and Benchmarking:** Event cameras have attracted considerable research interest due to their unique ability to capture rapid motion through asynchronous event processing [37], [38]. Recent studies [39]–[41] have begun to explore the potential of event data for high-level tasks, especially in the context of driving scenarios. Pioneering works, such as [17], [42], have collected large RGB-Event datasets tailored to urban scenes. Nevertheless, these datasets often lack comprehensive manual high-level annotations. Some approaches resort to off-the-shelf techniques to generate pseudo ground truth annotations, such as the use of YOLO bounding boxes [43] or semantic segmentation [44]. However, the quality of such pseudo annotations may be limited. Furthermore, these efforts typically focus on global detection without distinguishing between motion features, i.e., discerning moving objects from static ones. The concurrent works [45]–[47] detect moving objects but only a small number of sequences with indoor event input. In this paper, we introduce a novel large-scale MOS dataset that provides pixel-level annotations of scenes with moving ego vehicles. Our dataset stands as the first of its kind, designed specifically to tackle the distinctive challenges posed by moving objects in dynamic urban environments.

## III. METHODOLOGY

Given a video clip, denoted as  $\{I_t \in \mathbb{R}^{3 \times H \times W}\}_{t=1}^T$ , comprising a sequence of  $T$  frames, our objective is to predict masks  $\{P_t \in \mathbb{R}^{H \times W}\}_{t=1}^T$  corresponding to the moving objects within the frames. These predicted masks are intended to closely align with the ground truth masks  $\{M_t \in \mathbb{R}^{H \times W}\}_{t=1}^T$  for the respective frames.

In the following sections, we introduce our end-to-end learnable EmoFormer as shown in Figure 2. In contrast to the conventional approach of pre-computing the optical flow, our method harnesses motion cues directly from the recorded event data. Moreover, we seamlessly integrate the event prior as an additional dense supervision, resulting in event-free inference. This pivotal advance not only brings our model on par with RGB networks but also improves its suitability for real-world applications. For the sake of brevity in subsequent discussions, we omit the notion  $\{.,.\}_{t=1}^T$ .

### A. Prior Generation

Without loss of generalizability, we take the encoded RGB feature  $F \in \mathbb{R}^{c \times h \times w}$  from the last layer of the backbone from one image as an example. While the RGB feature is rich in appearance information, it inherently lacks awareness of temporal variations. Although inner-frame motion can be modeled by comparing two consecutive images, the video’s limited frames per second (FPS) restrict its ability to capture rapid motion. Hence, our objective is to estimate the rapid motion from a single RGB image, with the help of event data acquired during the frame time.

Specifically, we use a combination of depth-wise separable convolutions as the generation module. First, we project  $F$  into a deeper latent feature space using a  $Conv_{1 \times 1}$  operation, capture local motion awareness with  $DSC_{3 \times 3}$ , and then

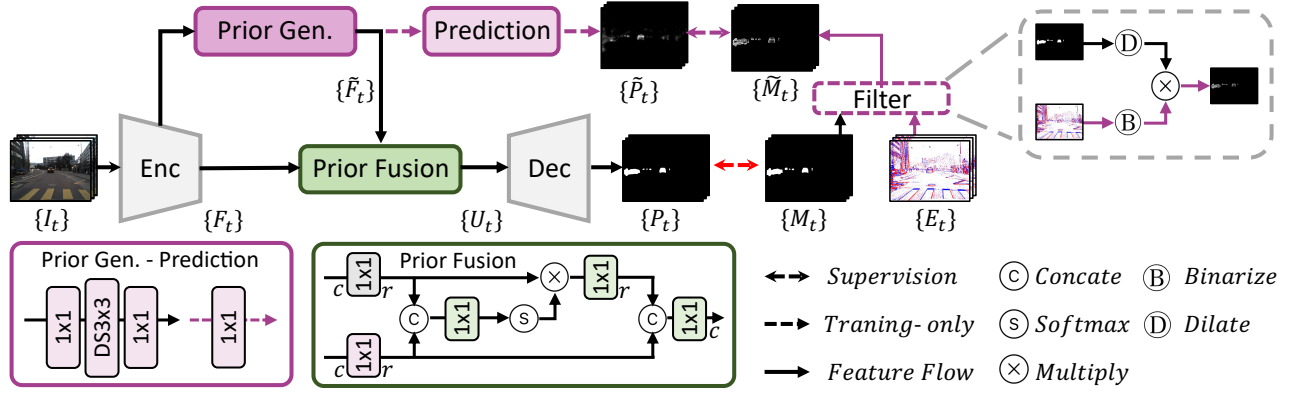


Fig. 2: **Architecture.** In addition to the standard RGB Encoder-Decoder architecture, we introduce an auxiliary branch dedicated to harnessing the motion insights derived from the recorded event data (Sec. III-A). This learned representation is subsequently merged into the main processing pipeline, thereby enhancing feature modeling (Sec. III-B). To further refine event-based learning and the understanding of object dynamics, we employ semantic maps to transfer global scene motion into targeting objects’ motion (Sec. III-C). Such a filter strategy leads to a tightly coupled semantic-guided event awareness, ultimately shaping our joint learning scheme.

project the result back using another  $Conv_{1 \times 1}$ . The choice of depth-wise separable convolution stems from the observation that rapid motion typically occurs within a localized region. Furthermore, the depth-wise convolution contributes to efficient processing.

Consequently, we obtain an intermediate representation corresponding to the high dynamic features, denoted as  $\tilde{F} \in \mathbb{R}^{c \times h \times w}$ . Mathematically, we have:

$$\tilde{F} = Conv_{1 \times 1}(DSCConv_{3 \times 3}(Conv_{1 \times 1}(F))). \quad (1)$$

To provide explicit guidance for the intermediate feature, we initially use a  $Conv_{1 \times 1}$  operation to generate a one-channel prediction map  $\tilde{P}$  based on  $\tilde{F}$ . We then integrate event supervision, as described in Sec. III-C.

### B. Prior Fusion

To seamlessly integrate the high dynamic feature  $\tilde{F}$  with the visual feature  $F$ , we introduce a novel prior fusion framework shown in Figure 2. Traditional RGB-Event fusion methods such as [43] often operate in an equal or higher dimensional feature space. However, the direct application of such methods to our context poses a challenge. Since one feature is estimated based on the other, even with event supervision, it may still exhibit a high degree of redundancy compared to the other. Consequently, these methods may retain redundant information rather than mutually complementary ones.

We propose to project both of these features into a smaller subspace, denoted as  $r \ll C$ . This projection can be likened to a low-rank decomposition, wherein only the most informative modality-specific features are preserved. Moreover, this approach naturally reduces the computational overhead compared to alternative fusion techniques. Mathematically, we achieve this projection as follows:

$$f = Conv_{1 \times 1}(F); \quad \tilde{f} = Conv_{1 \times 1}(\tilde{F}). \quad (2)$$

We then conduct cross-modal feature modeling within this lower-dimensional space  $r$ .

To facilitate this, we initially concatenate these feature maps, forming a correlation matrix  $f_c \in \mathbb{R}^{r \times h \times w}$ , followed by per-pixel attention using a global Softmax function:

$$Att = Softmax(f_c); \quad (3)$$

$$f_c = Conv_{1 \times 1}(Concat(f; \tilde{f})).$$

We then apply this attention map to the low-rank RGB feature, thereby mitigating the domain gap. Subsequently, we obtain the shared embedding by merging the attention-enhanced RGB feature and the motion feature by concatenation. Finally, this mixed embedding is projected back into the full-rank space  $C$  to yield the fused feature  $U \in \mathbb{R}^{c \times h \times w}$ :

$$U = Conv_{1 \times 1}(Concat(f \times Att; \tilde{f})). \quad (4)$$

### C. Semantic-Guided Event Filter

Event data provide invaluable insights into the global scene motion. However, the dense distribution of events is not solely limited to areas occupied by moving objects, but can also encompass the contours of static or background elements. This complexity arises from the fact that our ego vehicle is also in motion, which makes it difficult to discern the motion exclusively associated with the target objects.

To solve this dilemma, we propose a novel approach that combines event data with ground truth semantic supervision. This hybrid strategy allows us to pass from the overall scene motion into the object motion, constraining the network attention around the moving object with dense supervision.

Specifically, it involves suppressing background regions by element-wise multiplication with the spatial semantic mask. However, the direct multiplication of the spatial semantic mask with the temporal event map can inadvertently remove essential motion cues, especially when the event motion map extends beyond the object boundaries due to motion blur. To mitigate this, We employ a dilation operation on the

spatial mask with a specified kernel size before conducting semantic-guided event filter. Mathematical, let  $M \in \mathbb{R}^{H \times W}$  be the ground truth (GT) semantic mask,  $D \in \mathbb{R}^{3 \times 3}$  be the dilation matrix, and  $E \in \mathbb{R}^{H \times W}$  be the event data. The filtered map  $\tilde{M} \in \mathbb{R}^{H \times W}$  is obtained by:

$$\tilde{M} = (M \oplus D) \circ E, \quad (5)$$

where  $\oplus$  is the morphological dilation and  $\circ$  is the Hadamard product. In our application, we set  $D$  as an all-ones matrix, and we treat positive and negative polarization changes equally in the binary form of event data.

Subsequently, we utilize  $\tilde{M}$  to supervise the intermediate representation  $\tilde{P}$  introduced in Sec. III-A. To minimize the disparity between these two maps, we leverage conventional MSE loss as the auxiliary loss denoted as  $\tilde{\mathcal{L}}$ .

#### D. Overall Learning Pipeline

Our event-guided training methodology departs from the conventional RGB learning pipeline by introducing additional semantic-guided event filter. On top of our RGB baseline [20], we incorporate the extra prior generation module (Sec. III-A) and the prior fusion module (Sec. III-B). Our objective function can be summarized as follows:

$$\operatorname{argmin} (\mathcal{L}(\{P_t\}; \{M_t\}) + \tilde{\mathcal{L}}(\{\tilde{P}_t\}; \{\tilde{M}_t\})). \quad (6)$$

It is noteworthy that the event prediction  $\{\tilde{P}_t\}$  is exclusively employed during training to leverage event supervision, but it is not required during testing. During inference, our model operates solely on frames for outputting  $\{P_t\}$ , ensuring an event-free process.

### IV. OUR DSEC-MOS DATASET

#### A. Dataset Annotation

Our DSEC-MOS dataset builds upon the DSEC-MOD dataset [18], originally designed for moving object detection with bounding box annotations. To create DSEC-MOS, we followed a systematic process:

**Candidate Mask Generation:** We initiated the process by employing the modified SOTA segmentation model [48] with bounding box prompts to generate candidate masks.

**Annotator Training:** Next, we train three annotators to become proficient in video segmentation tasks. They are provided with guidance and instructions on well-established VOS benchmarks to familiarize themselves with the task.

**Manual Verification:** After training, each annotator was tasked with reviewing the candidate masks, aided by reference to adjacent frames. If a candidate mask was found to be suboptimal with a visually incomplete or inaccurate segmentation, the annotator was asked to correct it manually.

**Specialization by Illumination Condition:** Since our dataset comprises sequences captured under various illumination conditions, we assigned each annotator to sequences from one typical condition to develop specializations.

**Cross and Double Verification:** The sequences from each illumination condition underwent cross and double verification. In case of uncertainty or discrepancies, annotators were encouraged to discuss and jointly modify the masks until a consensus was reached.

TABLE I: Datasets Per-Attribute Comparison

Dataset	GT	Outdoor	Per-Frame	MOS
EV-IMO [45]	✓	✗	-	✓
DSEC-Sem [44]	✗	✓	✓	✗
DSEC Night-Sem [11]	✓	✓	✗	✗
<b>Ours</b>	✓	✓	✓	✓

#### B. Dataset Comparison

In Table I, we offer a comprehensive per-attribute comparison with established event-based segmentation datasets. Significantly, our dataset emerges as a pioneering initiative, distinguished as the first of its kind to provide dense MOS masks from a moving ego vehicle. While EV-IMO [45] stands as another notable MOS dataset, it is confined to indoor settings and lacks paired RGB frames. This absence hinders its applicability for assessing RGB models, underscoring the unique nature of our dataset.

## V. EXPERIMENTS

#### A. Experimental Setup

**Implementation Details:** We take video-swin [49] as the encoder that extracts features from 6 frames. Following [20], each input frame is resized to 360×409. Similar to [18], we leverage event data from 50ms, in accordance with the frame time. Classic data augmentation methods such as multi-scale training and flipping are adopted. We adopt AdamW as the optimizer, with an initial learning rate 1e-5 and 1e-4 for the backbone and others, respectively. The weight decay is set to 1e-4 and the polynomial learning rate decay is employed with a power of 0.9. Our network is built upon Pytorch. The whole network training takes around 2 days on a 3090 GPU. Unlike most VOS works, our method does not require any pretraining on large-scale datasets. For benchmarking, we retrain all the SOTA counterparts from their official resources. For practical application consideration, we exclude the time-burden CRF post-processing for all.

**Evaluation Metrics:** We use classical evaluation metrics, including region similarity  $\mathcal{J}$  and contour accuracy  $\mathcal{F}$ , in accordance with the VOS tasks [34]. We compute both the mean and recall measures for each metric, as well as the overall performance  $\mathcal{J}\&\mathcal{F}$  based on their mean measures.

#### B. Benchmark and Comparison

It is noteworthy that our benchmark does not include comparisons with RGB-Event object segmentation methods, as none have been identified thus far. Despite the availability of certain semantic segmentation techniques [11], [44], [50], [51], these methods do not deal with the distinction between moving and static objects, so they do not fall within the scope of our benchmark.

**Comparison against SOTA Methods:** Our moving object segmentation task is very similar to Automatic Video Object Segmentation (AVOS) tasks; in that we do not provide manual clues such as initial masks or bounding boxes to identify the target objects. Instead, we aim to automatically segment moving objects based solely on video sequences. To evaluate

TABLE II: **Quantitative Comparison** with SOTA AVOS Approaches on Our DSEC-MOS. Our EmoFormer outperforms both RGB and RGB-Flow counterparts with large margins. MS Inf. refers to Multi-Scale Inference. Our EmoFormer achieves the best performances during both single-scale and multi-scale inferences.

Method	Pub. & Year	Backbone	Optical Flow	MS Inf.	$\mathcal{J}$ Mean	$\mathcal{J}$ Recall	$\mathcal{F}$ Mean	$\mathcal{F}$ Recall	$\mathcal{J}$ & $\mathcal{F}$ Mean
RTNet [24]	CVPR'21	ResNet-101	✓	-	57.8	65.8	71.7	82.2	64.8
AMCNet [23]	ICCV'21	ResNet-101	✓	-	60.6	71.3	76.1	86.8	68.4
TMO [15]	WACV'23	ResNet-101	✓	-	56.2	62.1	73.0	82.8	64.6
Isomer [19]	ICCV'23	Swin-T	✓	-	55.9	65.3	72.2	86.0	64.1
F2Net [21]	AAAI'21	ResNet-50	-	-	45.1	48.4	66.9	74.5	56.0
TMO [15]	WACV'23	ResNet-101	-	-	53.5	57.1	70.3	81.8	61.9
MED-VT [20]	CVPR'23	Swin-B	-	-	58.0	67.3	81.0	90.6	69.5
<b>EmoFormer</b>	<b>Ours</b>	<b>Swin-B</b>	<b>-</b>	<b>-</b>	<b>66.5</b>	<b>72.9</b>	<b>82.2</b>	<b>92.2</b>	<b>74.4</b>
HFAN [14]	ECCV'22	Swin-T	✓	✓	62.6	69.5	76.8	85.7	69.7
MED-VT [20]	CVPR'23	Swin-B	-	✓	63.5	71.8	81.5	90.3	72.5
<b>EmoFormer</b>	<b>Ours</b>	<b>Swin-B</b>	<b>-</b>	<b>✓</b>	<b>68.9</b>	<b>77.7</b>	<b>83.7</b>	<b>93.1</b>	<b>76.3</b>




Fig. 3: **Qualitative Comparison.** Our generated masks are closer to the Ground Truth (GT) compared to the counterparts. Please zoom in for details.

the effectiveness of our method, we compared 8 state-of-the-art automatic VOS methods as reported in Table II. To ensure fair comparisons, we evaluated our performance with both single-scale and multi-scale inference techniques. The results demonstrate a significant performance improvement over all other methods. In the case of single-scale inference, our network significantly improves the performance of our RGB baseline (MED-VT) with an absolute gain of +8.5%  $\mathcal{J}$  Mean, outperforming the current state-of-the-art method, AMCNet [23]. It is worth noting that AMCNet requires both RGB and flow maps during inference, whereas our method operates only on RGB frames. In the case of multi-scale inference, we achieve an absolute gain of +5.4%  $\mathcal{J}$  Mean over MED-VT. We also provide qualitative comparisons in Figure 3, illustrating that our method excels at exploring temporal cues and generating prediction masks that closely match ground truth, outperforming any pure RGB or RGB-Flow counterparts.

### C. Ablation Studies

We perform multi-scale inference for ablation studies on Supervision Strategies. In our training pipeline, we leverage event data as auxiliary supervision, facilitated by a semantic mask. To assess the effectiveness of this approach, we conducted experiments in which we replaced the intermediate representation with alternative supervision methods, such as optical flow only, semantic mask only, and event data only. The quantitative results are shown in Table III. It is evident

TABLE III: **Ablation Study on Supervision Strategies.**  $\circ$  is the Hadamard multiplication to suppress static responses.

Aux. Sup. Source	$\circ$ GT	Dilation	$\mathcal{J}$ M	$\mathcal{F}$ M	$\mathcal{J}$ & $\mathcal{F}$ M
1	-	-	63.5	81.5	72.5
2	Flow	-	66.2	80.7	73.5
3	Semantic	-	67.2	82.2	74.7
4	Semantic	✓	67.6	83.1	75.4
5	Event	-	65.4	80.1	72.8
6	Event	✓	67.9	82.8	75.4
7	Event	✓	<b>68.9</b>	<b>83.7</b>	<b>76.3</b>

that, while the auxiliary supervisions consistently improve our baseline, their individual performance is inferior to our full approach. We observed that omitting the semantic map guidance and the dilation operation leads to performance deterioration. This outcome can be attributed to the role of semantic guidance in attenuating noisy event responses from static regions. In addition, dilation plays a crucial role in the alignment of spatial masks with temporal motion blur.

## VI. CONCLUSION

We demonstrate the effectiveness of using the event prior for moving object segmentation in dynamic scenes. Our method combines the inherent advantages of event data for capturing global motion with semantic supervision. This leads to an intermediate representation that enables the network to separate moving objects from the static background, resulting in another level of dense supervision. Moreover, our method no longer requires any event input during inference, which makes it particularly efficient and suitable for real-world applications based on standard RGB cameras. We also introduce a new large-scale dataset designed for moving object segmentation in dynamic autonomous driving scenarios, which is the first of its kind. Our method significantly outperforms the state of the art and confirms the potential of event-based pipelines for autonomous driving. To sum up, we hope that our method and dataset can promote further research in this field.

## REFERENCES

- [1] Y. Wu, R. Gao, J. Park, and Q. Chen, "Future video synthesis with object motion prediction," in *CVPR*, 2020.

- [2] M. Tancik, V. Casser, X. Yan, S. Pradhan, B. Mildenhall, P. P. Srinivasan, J. T. Barron, and H. Kretzschmar, "Block-nerf: Scalable large scene neural view synthesis," in *CVPR*, 2022.
- [3] A. Kundu, K. Genova, X. Yin, A. Fathi, C. Pantofaru, L. J. Guibas, A. Tagliasacchi, F. Dellaert, and T. Funkhouser, "Panoptic neural fields: A semantic object-aware neural scene representation," in *CVPR*, 2022.
- [4] J. Xie, W. Xie, and A. Zisserman, "Segmenting moving objects via an object-centric layered representation," *NeurIPS*, 2022.
- [5] A. Philipp and D. Goehring, "Analytic collision risk calculation for autonomous vehicle navigation," in *ICRA*. IEEE, 2019.
- [6] Y. Liang, J. Liu, D. Zhang, and Y. Fu, "Mpi-flow: Learning realistic optical flow with multiplane images," in *ICCV*, 2023.
- [7] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, et al., "Event-based vision: A survey," *IEEE TPAMI*, vol. 44, no. 1, pp. 154–180, 2020.
- [8] S. Guo and T. Delbruck, "Low cost and latency event camera background activity denoising," *IEEE TPAMI*, vol. 45, no. 1, pp. 785–795, 2022.
- [9] G. Chen, H. Cao, J. Conradt, H. Tang, F. Rohrbein, and A. Knoll, "Event-based neuromorphic vision for autonomous driving: A paradigm shift for bio-inspired visual sensing and perception," *IEEE SPM*, vol. 37, no. 4, pp. 34–49, 2020.
- [10] A. I. Maqueda, A. Loquercio, G. Gallego, N. García, and D. Scaramuzza, "Event-based vision meets deep learning on steering prediction for self-driving cars," in *CVPR*, 2018.
- [11] R. Xia, C. Zhao, M. Zheng, Z. Wu, Q. Sun, and Y. Tang, "Cmnda: Cross-modality domain adaptation for nighttime semantic segmentation," in *ICCV*, 2023.
- [12] T. Zhou, F. Porikli, D. J. Crandall, L. Van Gool, and W. Wang, "A survey on deep learning technique for video segmentation," *IEEE TPAMI*, vol. 45, no. 6, pp. 7099–7122, 2022.
- [13] Y.-T. Hu, J.-B. Huang, and A. G. Schwing, "Unsupervised video object segmentation using motion saliency-guided spatio-temporal propagation," in *ECCV*, 2018.
- [14] G. Pei, F. Shen, Y. Yao, G.-S. Xie, Z. Tang, and J. Tang, "Hierarchical feature alignment network for unsupervised video object segmentation," in *ECCV*. Springer, 2022.
- [15] S. Cho, M. Lee, S. Lee, C. Park, D. Kim, and S. Lee, "Treating motion as option to reduce motion dependency in unsupervised video object segmentation," in *WACV*, 2023.
- [16] K. Zhang, Z. Zhao, D. Liu, Q. Liu, and B. Liu, "Deep transport network for unsupervised video object segmentation," in *ICCV*, 2021.
- [17] M. Gehrig, W. Aarents, D. Gehrig, and D. Scaramuzza, "Dsec: A stereo event camera dataset for driving scenarios," *IEEE RAL*, vol. 6, no. 3, pp. 4947–4954, 2021.
- [18] Z. Zhou, Z. Wu, R. Bouteau, F. Yang, C. Demonceaux, and D. Ginjac, "Rgb-event fusion for moving object detection in autonomous driving," in *ICRA*. IEEE, 2023.
- [19] Y. Yuan, Y. Wang, L. Wang, X. Zhao, H. Lu, Y. Wang, W. Su, and L. Zhang, "Isomer: Isomeric transformer for zero-shot video object segmentation," in *ICCV*, 2023.
- [20] R. Karim, H. Zhao, R. P. Wildes, and M. Siam, "Med-vt: Multiscale encoder-decoder video transformer with application to object segmentation," in *CVPR*, 2023.
- [21] D. Liu, D. Yu, C. Wang, and P. Zhou, "F2net: Learning to focus on the foreground for unsupervised video object segmentation," in *AAAI*, 2021.
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *NeurIPS*, vol. 30, 2017.
- [23] S. Yang, L. Zhang, J. Qi, H. Lu, S. Wang, and X. Zhang, "Learning motion-appearance co-attention for zero-shot video object segmentation," in *ICCV*, 2021.
- [24] S. Ren, W. Liu, Y. Liu, H. Chen, G. Han, and S. He, "Reciprocal transformations for unsupervised video object segmentation," in *CVPR*, 2021.
- [25] W. Wang, H. Song, S. Zhao, J. Shen, S. Zhao, S. C. Hoi, and H. Ling, "Learning unsupervised video object segmentation through visual attention," in *CVPR*, 2019.
- [26] Y. Yang, B. Lai, and S. Soatto, "Dystab: Unsupervised object segmentation via dynamic-static bootstrapping," in *CVPR*, 2021.
- [27] G. Sun, Y. Liu, H. Tang, A. Chhatkuli, L. Zhang, and L. , "Mining relations among cross-frame affinities for video semantic segmentation," in *ECCV*. Springer, 2022.
- [28] Y. Zhou, X. Xu, F. Shen, X. Zhu, and H. T. Shen, "Flow-edge guided unsupervised video object segmentation," *IEEE TCSVT*, vol. 32, no. 12, pp. 8116–8127, 2021.
- [29] M. Cho, T. Kim, W. J. Kim, S. Cho, and S. Lee, "Unsupervised video anomaly detection via normalizing flows with implicit latent features," *PR*, vol. 129, p. 108703, 2022.
- [30] Y. Yang, A. Loquercio, D. Scaramuzza, and S. Soatto, "Unsupervised moving object detection via contextual information separation," in *CVPR*, 2019.
- [31] M. Sultana, A. Mahmood, and S. K. Jung, "Unsupervised moving object detection in complex scenes using adversarial regularizations," *IEEE TMM*, vol. 23, pp. 2005–2018, 2020.
- [32] T. Zhou, S. Wang, Y. Zhou, Y. Yao, J. Li, and L. Shao, "Motion-attentive transition for zero-shot video object segmentation," in *AAAI*, 2020.
- [33] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool, "The 2017 davis challenge on video object segmentation," *arXiv preprint arXiv:1704.00675*, 2017.
- [34] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung, "A benchmark dataset and evaluation methodology for video object segmentation," in *CVPR*, 2016.
- [35] S. Caelles, A. Montes, K.-K. Maninis, Y. Chen, L. Van Gool, F. Perazzi, and J. Pont-Tuset, "The 2018 davis challenge on video object segmentation," *arXiv preprint arXiv:1803.00557*, 2018.
- [36] H. Ding, C. Liu, S. He, X. Jiang, P. H. Torr, and S. Bai, "Mose: A new dataset for video object segmentation in complex scenes," in *ICCV*, 2023.
- [37] Y.-F. Zuo, J. Yang, J. Chen, X. Wang, Y. Wang, and L. Kneip, "Devo: Depth-event camera visual odometry in challenging conditions," in *ICRA*. IEEE, 2022.
- [38] R. W. Baldwin, R. Liu, M. Almatrafi, V. Asari, and K. Hirakawa, "Time-ordered recent event (tore) volumes for event cameras," *IEEE TPAMI*, vol. 45, no. 2, pp. 2519–2532, 2022.
- [39] A. Amini, T.-H. Wang, I. Gilitschenski, W. Schwarting, Z. Liu, S. Han, S. Karaman, and D. Rus, "Vista 2.0: An open, data-driven simulator for multimodal sensing and policy learning for autonomous vehicles," in *ICRA*. IEEE, 2022.
- [40] K. Ta, D. Bruggemann, T. Brödermann, C. Sakaridis, and L. Van Gool, "L2e: Lasers to events for 6-dof extrinsic calibration of lidars and event cameras," in *ICRA*. IEEE, 2023.
- [41] Z. Ding, R. Zhao, J. Zhang, T. Gao, R. Xiong, Z. Yu, and T. Huang, "Spatio-temporal recurrent networks for event-based optical flow estimation," in *AAAI*, 2022.
- [42] A. Z. Zhu, D. Thakur, T. Özslan, B. Pfrommer, V. Kumar, and K. Daniilidis, "The multivehicle stereo event camera dataset: An event camera dataset for 3d perception," *IEEE RAL*, vol. 3, no. 3, pp. 2032–2039, 2018.
- [43] A. Tomy, A. Paigwar, K. S. Mann, A. Renzaglia, and C. Laugier, "Fusing event-based and rgb camera for robust object detection in adverse conditions," in *ICRA*. IEEE, 2022.
- [44] Z. Sun, N. Messikommer, D. Gehrig, and D. Scaramuzza, "Ess: Learning event-based semantic segmentation from still images," in *ECCV*. Springer, 2022.
- [45] A. Mitrokhin, C. Ye, C. Fermüller, Y. Aloimonos, and T. Delbruck, "Ev-imo: Motion segmentation dataset and learning pipeline for event cameras," in *IROS*. IEEE, 2019.
- [46] S. Zhang, L. Sun, and K. Wang, "A multi-scale recurrent framework for motion segmentation with event camera," *IEEE Access*, 2023.
- [47] Y. Zhou, G. Gallego, X. Lu, S. Liu, and S. Shen, "Event-based motion segmentation with spatio-temporal graph cuts," *IEEE TNLS*, 2021.
- [48] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al., "Segment anything," in *ICCV*, 2023.
- [49] Z. Liu, J. Ning, Y. Cao, Y. Wei, Z. Zhang, S. Lin, and H. Hu, "Video swin transformer," in *CVPR*, 2022.
- [50] Z. Jia, K. You, W. He, Y. Tian, Y. Feng, Y. Wang, X. Jia, Y. Lou, J. Zhang, G. Li, et al., "Event-based semantic segmentation with posterior attention," *IEEE TIP*, vol. 32, pp. 1829–1842, 2023.
- [51] J. Zhang, K. Yang, and R. Stiefelhagen, "Issafe: Improving semantic segmentation in accidents by fusing event-based data," in *IROS*. IEEE, 2021.