

Cross-modal Orthogonal High-rank Augmentation for RGB-Event Transformer-trackers

Zhiyu Zhu, Junhui Hou*, and Dapeng Oliver Wu

Department of Computer Science, City University of Hong Kong

zhiyuzhu2-c@my.cityu.edu.hk; jh.hou@cityu.edu.hk; dapengwu@cityu.edu.hk

Abstract

This paper addresses the problem of cross-modal object tracking from RGB videos and event data. Rather than constructing a complex cross-modal fusion network, we explore the great potential of a pre-trained vision Transformer (ViT). Particularly, we delicately investigate plug-and-play training augmentations that encourage the ViT to bridge the vast distribution gap between the two modalities, enabling comprehensive cross-modal information interaction and thus enhancing its ability. Specifically, we propose a mask modeling strategy that randomly masks a specific modality of some tokens to enforce the interaction between tokens from different modalities interacting proactively. To mitigate network oscillations resulting from the masking strategy and further amplify its positive effect, we then theoretically propose an orthogonal high-rank loss to regularize the attention matrix. Extensive experiments demonstrate that our plug-and-play training augmentation techniques can significantly boost state-of-the-art one-stream and two-stream trackers to a large extent in terms of both tracking precision and success rate. Our new perspective and findings will potentially bring insights to the field of leveraging powerful pre-trained ViTs to model cross-modal data. The code is publicly available at <https://github.com/ZHU-Zhiyu/High-Rank-Event-Tracker>.

1. Introduction

Event cameras asynchronously capture pixel intensity fluctuations with an ultra-high temporal resolution, low latency, and high dynamic range, making it gain increasing attention recently [38, 42, 15]. Owing to such admirable advantages, event cameras have been widely adopted in various applications, such as object detection [38, 30, 39, 42, 11] and depth/optical flow estimation [16, 63]. Particularly,

*Corresponding author: Junhui Hou. This work was supported in part by the Hong Kong Research Grants Council under Grant 11218121 and Grant 11202320, and in part by the Hong Kong Innovation and Technology Fund under Grant MHP/117/21.

the distinctive sensing mechanism makes event cameras to be a promising choice for object tracking [45, 26, 60, 64, 18, 19].

Despite many advantages of event-based object tracking under special environments, e.g., low-light, high-speed motion, and over-exposed, event data lack sufficient visual cues, such as color, texture, and complete contextual appearance that can be easily captured by RGB data, resulting in only event-based vision still suffering from relatively inferior performance in practice. Thus, a more promising direction is to investigate cross-modal object tracking from both RGB and event data, where the merits of the two modalities can be well leveraged for pursuing higher performance. However, the vast distribution gap between RGB and event data poses significant challenges in designing algorithms for modeling cross-modal information. Most existing pioneering cross-modal trackers heavily engage in robust cross-modal fusion modules, which is cumbersome to use advanced embedding backbones for boosting performance.

In view of the success of Transformer-based tracking algorithms [31, 59, 54, 7, 62], where the multi-head attention naturally models the indispensable correlation relationship between template and search regions, we plan to investigate the potential of pre-trained powerful vision Transformers (ViTs) in cross-modal object tracking from both RGB and event data. However, those pre-trained Transformers with RGB data may not be able to fully model the *essential* feature interaction across RGB and event data, due to the distribution gap between the two modalities. To this end, we study *plug-and-play* training techniques for augmenting the pre-trained Transformer used as the embedding backbone of our RGB-event object tracking framework.

To be specific, to promote the learning of the attention layer across two modalities, we propose a cross-modal mask modeling strategy, which randomly masks/pops out the multi-modal tokens. We anticipate that, in reaction to the absence of a particular modality at certain locations, the network would proactively enhance interactions on the remaining cross-modal tokens. Nevertheless, randomly masking

tokens will inevitably alter data distributions and introduce disruptions, impeding network training. To mitigate the induced negative effect, we further propose a regularization term to guide the training of each attention layer. Based on the observation that the values of internal attention matrices of a Transformer indicate the degree of cross-modal feature interaction, we propose to orthogonalize the attention matrix to promote its rank obligatorily. Beyond, we anticipate that such regularization could encourage the cross-modal correlation to be evenly and concisely established using the multi-domain signatures, rather than unduly reliant on a specific domain. Finally, we apply the proposed techniques to state-of-the-art one-stream and two-stream Transformer-based tracking frameworks and experimentally demonstrate that their tracking performance is further boosted significantly.

In summary, the contributions of this paper are:

- a mask modeling strategy for encouraging the interaction between the cross-modal tokens in a *proactive* manner;
- *theoretical* orthogonal high-rank regularization for suppressing network fluctuations induced by cross-modal masking while amplifying its positive effect; and
- new state-of-the-art baselines for RGB-event object tracking.

Last but not least, our novel perspectives will potentially bring insights to the field of leveraging pre-trained powerful ViTs to process and analyze cross-modal data.

2. Related Work

2.1. Object Tracking

Recent years have seen remarkable progress in the study of object tracking, which is primarily due to the widespread success of deep learning [1, 27]. Based on the distribution of computational burdens, current methods could be generally divided into two-stream [1, 2, 31] and one-stream methods [57, 6]. As the earlier invented and relatively mature ones, most offline Siamese-based tracking methods [1, 28, 27] fall into the first category. It utilizes a delicate embedding backbone to extract semantic-rich embeddings and then models the target location via either a direct proposal head [1] or an online optimization process [2], which is also called deep Siamese-trackers or discriminative correlation filters, respectively [23]. SiamFC [1] first developed a fully-convolutional architecture to fuse template and search embeddings for object tracking. Though introducing a single-stage RPN [46] detector SiamRPN [28] achieved target object tracking by comparing the current-frame features to those from a template. To remove the disturbance factors, e.g., padding, SiamRPN++ [27] introduced a spatial-aware sampling strategy and further utilized

ResNet [21] to embed representative features for Siamese-based tracking. DiMP [2] proposed to exploit both target and background appearances to achieve object tracking. KYS [3] represented the scene information as dense state vectors and utilizes such state vectors to maximize the tracking performance. Besides, some spatio-temporal-based methods also exploit temporal information to achieve robust and effective tracking [40, 24, 53, 56]. MDNet [40] separated domain-independent from domain-specific information via a CNN-based framework. RT-MDNet [24] further improved it via an RoI-Align strategy, which extracts more precise embeddings from feature maps of targets and candidates. Swin-tracker [31] introduced the Swin-Transformer [34] to effectively encode the semantic information from input images for high-performance visual tracking.

Due to the extraordinary correlation modeling ability of Transformer, an emerging branch of one-stream methods shows strong potential in correlation modeling. OS-track [57] unified the embedding and relation modeling processes with a single vanilla ViT [14], which achieves admirable performance with reduced computational resources. Meanwhile, SimViT-Track [6] proposed a similar approach, which feeds search and template image tokens straight into a ViT backbone and performs regression and classification on the resulting tokens.

In summary, with the success of existing embedding backbones, such as ViT [14] and Swin-Transformer [34], more intriguing and effective methods have been proposed recently. While these methods could achieve admirable performance, most of them are driven by matching semantically identical segments of the search and template regions viewed as RGB images. As a result, their performance is inextricably tied to imaging characteristics, which can be compromised in specific scenarios such as high-speed and low-light scenes. Hence, it is highly desired to incorporate multi-modal inputs to remedy each deficiency. Moreover, the crucially multi-modal data necessitates additional efforts to generalize these methods to the event-based.

2.2. Event-based Tracking

Owing to its innate characteristics and superiority for object tracking, event-based tracking has been a progressively prevalent subject for research in recent years. Additionally, existing approaches may be broadly classified into two categories: model-based and data-driven. Through describing surrounding environments by a photometric 3D map, Bryner *et al.* [4] proposed to track the 6-DOF pose of a camera. To capture the spatio-temporal geometry of event data, Mitrokhin *et al.* [38] utilized a parametric model to compensate camera motion. Based on a pipeline of tracking-learning-detection, Ramesh *et al.* [44] proposed an object tracking algorithm for event cameras, which is

the first learning-based long-term event tracker. Then, Li *et al.* [29] introduced the VGG-Net-16 to encode the appearance of the event-stream object. Inspired by the classic Siamese-matching paradigm, Chae *et al.* [5] presented to track objects via learning an edge-aware similarity in the event domain. Recently, Zhang *et al.* [64], introduced a spiking transformer for encoding spatio-temporal information of object tracking. Moreover, ZHU *et al.* [64] proposed to utilize inherent motion information of event data to achieve effective object tracking. To summarize, although there are some promising studies that provide directive insights for event-based tracking, a limited number of works have sought to find complementary information from RGB data, e.g., semantic information.

2.3. Cross-modal Learning

Fusing embedding with multiple modalities is a sensible solution for perceiving and recognizing the objects robustly and accurately [43, 61, 65]. However, for current machine learning algorithms, learning representative patterns from multiple modalities is still a challenging issue [32, 22]. Wang *et al.* [51] proposed to apply data augmentation techniques to boost cross-modal 3D object detection. Liu *et al.* [32] utilized cross-modal feature rectification and fusion models for image segmentation with input from multiple modalities. Jaritz *et al.* [22] solved the multi-modal segmentation issue from the perspective of unsupervised domain adaptation. Moreover, Wang *et al.* [52] designed an RGB-T tracking framework by propagating the intermodal pattern and long-term context. Ye *et al.* [58] proposed a cross-modal self-attention module to achieve natural language-based image segmentation via adaptively capturing informative words and important regions in images. Zeng *et al.* [52] proposed to project the camera features onto the point set on LiDAR. In summary, recent works are clearly founded on network architecture, as is evident by their prevalence. Moreover, the current advanced Transformer paradigm could adaptively process different modalities. However, there is still a lack of further investigations and analysis of the internal mechanism.

3. Proposed Method

3.1. Motivation

Learning the correlation between the template and search regions robustly and precisely is one of the most essential aspects of object tracking. Fortunately, with current advancements in the multi-head attention mechanism, such correlation could be naturally achieved via Transformer-based frameworks [54, 7]. However, current powerful ViTs were usually pre-trained with RGB data, e.g., ImageNet [12], potentially resulting in that they cannot be adequately adapted to cross-modal learning, i.e., the full fea-

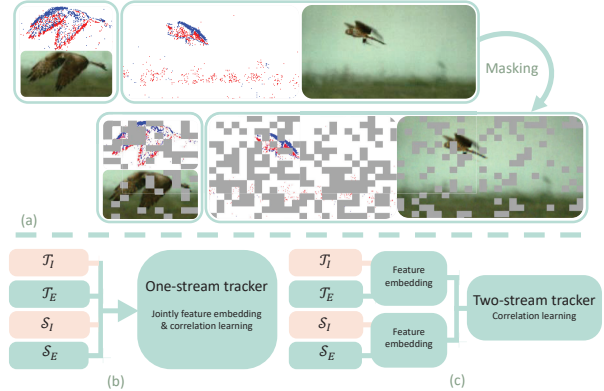


Figure 1. Illustration of (a) the proposed cross-modal mask modeling strategy on template and search data. General structures of Transformer-based RGB-event trackers (b) one-stream and (c) two-stream, where \mathcal{T} and \mathcal{S} represent the tokens of template and search patches, with subscripts I and E indicating the RGB and event modalities, respectively.

ture interaction between RGB and event data cannot be well achieved, which is essential for cross-modal object tracking, due to the vast distribution gap between RGB and event data. Accordingly, the tracking performance may be limited.

Instead of following existing cross-modal research paradigms mainly focused on designing sophisticated cross-modal information fusion networks, we aim to explore *plug-and-play training augmentation techniques* to mitigate the above-mentioned potential limitation of a pre-trained ViT used as the embedding backbone of an RGB-Event object tracking scheme. Generally, based on a fundamental and essential premise that different modalities possess their own unique benefits for a cross-modal tracker, token embedding information should be adequately transmitted across multimodalities, especially for the regions with target objects, in a bid to enhance themselves using specific merits from the other modality. Thus, we propose a mask modeling strategy to enable the network to proactively exploit the cross-modal information in Sec. 3.2. Furthermore, we propose a high-rank orthogonalization mechanism in Sec. 3.3, which can not only alleviate network fluctuations induced by the mask modeling strategy but also further boost cross-modal information interaction.

In what follows, we will detail the proposed techniques adapted to both one-stream and two-stream trackers, as illustrated in Fig. 1 (b) and Fig. 1 (c), respectively. We always use I and E in the subscripts to indicate the RGB and event modalities, and \mathcal{T} and \mathcal{S} are the tokens of template and search regions, respectively.

3.2. Mask-driven Cross-modal Interaction

Grouping tokens via similarity is one of the most representative steps for the self-attention mechanism of a Trans-




Figure 2. Illustration of the underlying intuition of our cross-modal strategy, where s , and t denote tokens from search and template regions, respectively. The width of the lines indicates the degree of cross-modal interaction, i.e., the thicker, the more comprehensive. (a) The correlation between tokens is insignificant in the baseline model. (b) During training, tokens are *randomly* masked to facilitate cross-modal interaction. The random manner ensures the potential token interaction routes are strengthened with an equal probability. (c) The augmented model can cross the gap between cross-modal tokens to interact with each other.

former [49]. However, due to the distribution gap between tokens corresponding to different modalities, the similarity-driven attention may tend to aggregate information from the identical modality, hence impeding the cross-modal learning. Thus, how to effectively and efficiently promote the cross-modal interactions is *critical* for maximizing the potential of a pre-trained ViT for RGB-event object tracking.

We propose a cross-modal mask modeling strategy to address this issue in a *proactive* manner, shown as Fig. 1 (a). As illustrated in Fig. 2, the underlying intuition of this strategy is through removing the patches of different modalities and locations, we expect that the task loss would enforce the network to spontaneously enhance/build cross-modal correlation, due to the remaining tokens in different modalities. Once the interaction is established, the RGB and event tokens may learn to shrink the distribution gap, maintaining such correlation to the inference phase. Specifically, we apply random masks to RGB and event data to remove distinct patches. To begin, for the one-stream methods, masking elements can be readily accomplished by simply popping out corresponding elements, which could concurrently lessen the network training burden. For the two-stream methods, due to the large computational resource consumption of the embedding backbone, we directly average the masked features of RGB and event data at the primary stage, which are further fed into the high-level embedding backbone and relation modeling modules for the object proposal.

Remark. It is worth noting that the motivation and objective of the proposed masking strategy are considerably **different** from those of the well-known masked image modeling [20, 55, 25]. We start from the pursuit of promoting the network to actively utilize cross-modal information. Thus, the patches with distinct positions across RGB and event modalities are randomly removed to permit each location

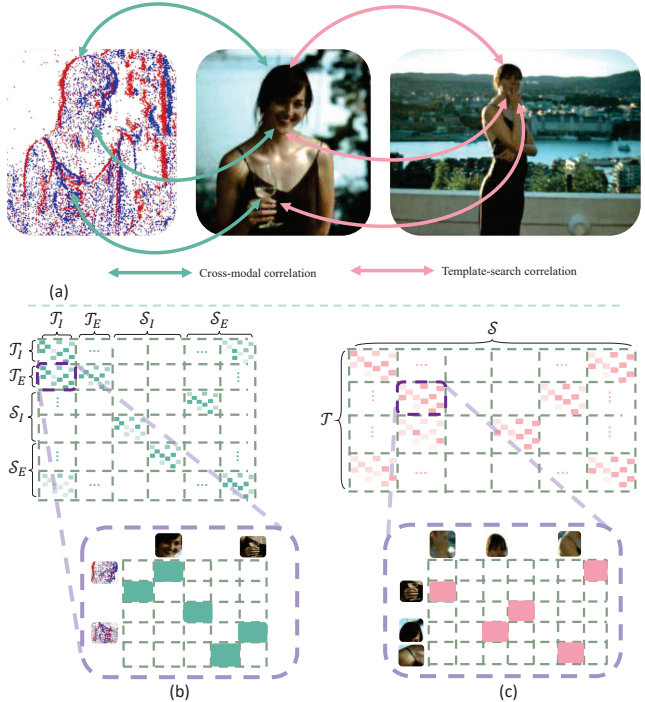


Figure 3. (a) Illustration of the correlation modeling for the cross-modal tracker, where **left arrows** reflect the desirable *cross-modal correlation* and **right ones** indicate the satisfactory *search-template correlation*. The matrix form of such correlations is shown (b) and (c). Each row consists of attention values of a **query** to multiple **keys**, i.e., the summation of each row is equal to 1. Here, we use the indicators of the row and column to represent the corresponding block, e.g., M_{T_E, T_I} indicates the matrix block located at the 2^{nd} row and the 1^{st} column. We also show the real matrices in Fig. 7.

can be perceived by the network but with different modalities. However, mask image modeling pre-trains network weights to comprehend image semantics by feeding just a subset of image patches to reconstruct the unseen area.

Although such a masking strategy used in the training phase is expected to strengthen the ability of the network to perceive cross-modal information to some extent, the randomly dropped information would potentially result in an unstable training process. Moreover, such disruptions are especially devastating for one-stream algorithms, which must concurrently learn representative embeddings and establish the relationship between the cross-modal template and search tokens (see the experimental demonstration in Sec. 4.2). Thus, to pull the network out of this predicament, we further propose orthogonal high-rank regularization in a *theoretical* manner in the next section.

3.3. Orthogonal High-rank Regularization

To appreciate the multi-head attention mechanism, we take a one-stream tracker [47] with the vanilla ViT [14] as

an example. As illustrated in Fig. 3 (b), its internal self-attention layers concurrently perceive the RGB and event tokens from both the template and search areas. Depending on the query and key belongings $k \in \mathbb{R}$, we can partition the resulting attention matrix into k^2 blocks (Here $k = 4$). Note that the attention values of a typical block reflect the degree of the interaction between tokens.

To mitigate network disturbs induced by the cross-modal mask modeling strategy and further amplify its positive effect (i.e., boosting cross-modal learning), we concentrate on the cross-modal zones of the attention matrix, such as M_{S_I, S_E} , and M_{S_E, S_I} . Assuming that if tokens are well-embedded and with highly discriminative features, each token will form a unique correlation with its identical counterpart, resulting in each row or column being orthogonal to the others. Moreover, as attention elements are non-negative, the corresponding matrix should be *full rank*¹. Therefore, we propose the following regularization to encourage some desired blocks of the attention matrix to be high-rank:

$$L(M, \tau) = \|\text{diag}(\Sigma) - \text{vec}(\tau)\|_1, M = U\Sigma V, \quad (1)$$

where $\tau \in \mathbb{R}$ is a pre-defined threshold value, $U \in \mathbb{R}^{n \times n}$, $\Sigma \in \mathbb{R}^{n \times m}$, and $V \in \mathbb{R}^{m \times m}$ are the outputs of the singular value decomposition (SVD) of block $M \in \mathbb{R}^{n \times m}$, and $\text{diag}(\cdot)$ returns a vector, consisting of the main diagonal elements of the input matrix, and $\text{vec}(\cdot)$ converts an input scalar to be a vector by duplicating the scalar. We impose the regularization term onto a set of blocks of the attention matrix $\{M^{(i)}\}_{i=1}^N$ standing for the interaction of cross-modal tokens. Due to its strong regularization effect, we empirically select the blocks corresponding to image-to-event attention (i.e. M_{S_I, \mathcal{T}_E} , and M_{S_I, S_E}), and the blocks to event-to-image attention (i.e., M_{S_E, \mathcal{T}_I} , and M_{S_E, S_I}). Moreover, as computing the SVD of a matrix is time-consuming, we randomly choose a layer to implement this regularization at each optimization step, instead of operating it in each layer.

For the two-stream methods, since the input data from different modalities are mixed in a preceding embedding backbone as shown in Fig. 1 (c), e.g., swin-Transformer [34], the resulting attention matrix only consists of two parts, i.e., the search-to-template and template-to-search regions, as illustrated in Fig. 3 (c). Under this scenario, we anticipate that the discriminative cross-modal tokens will be able to form a unique correlation with the identical object parts across template and search areas. As shown in the right part of Fig. 3 (a) and Fig. 3 (c), such a relationship would also produce that each row is orthogonal to the others. Thus, we also regularize the regions belonging to the target objects in $M_{S, \mathcal{T}}$. Specifically, guided by bounding box information, we first mask the attention weights in

non-target regions of $M_{S, \mathcal{T}}$, then apply Eq. (1) to increase the rank of the masked matrix.

3.4. Training

To train a Transformer-based tracker with the proposed plug-and-play augmentation techniques, at each optimization step, we first randomly mask/pop out event and image patches with a ratio of δ_e and δ_i ($0 < \delta < 1$), respectively. Then, we train the whole network with the following loss function:

$$L_{all} = L_{task} + \alpha L(M, \tau), \quad (2)$$

where L_{task} denotes the original task loss function, composed of regression and classification branches, and α is a balanced weight for the proposed regularization term.

4. Experiment

Implementation details. We evaluated the proposed plug-and-play training augmentation techniques on both one-stream and two-stream trackers. We set template and search sizes as 128 and 256, respectively, which contain $2 \times$ and $4 \times$ regions than annotations. Moreover, the location and scale jitter factors of the search region are set as 3 and 0.25, respectively (No jitter to template region). For one-stream, we directly adopted the SOTA method named color-event unified tracking (CEUTrack) [47] as our baseline model (ViT-B). Moreover, we also adapted ViT-L to CEUTrack to further validate the effectiveness of the proposed regularization term. During training, we used the same optimizer (ADAW), learning rate scheduler, and task loss function as the original paper. We set the batch size as 24 and the augmentation weight α in Eq. (2) empirically as 1.2. The masking ratios of both modalities δ_i and δ_e were set to 0.1.

For two-stream, to the best of our knowledge, there is no Transformer-based RGB-event tracker available, we chose the most recent event cloud-based motion-aware tracker (MonTrack) [64] and modified it with the proposal head of a Transformer-tracker [7] and the backbone of pre-trained swin-v2 [33] to construct two-stream RGB-event trackers for the detailed architecture). Moreover, we tested lightweight and heavy backbones, i.e., Swin-V2-Tiny [33] and Swin-V2-Base [33], to achieve comprehensive evaluation, and the resulting baselines are named MonTrack-T and MonTrack-B, respectively. To train the whole framework, we utilized the AdamW optimizer [35] with the learning rate of $1e^{-4}$ for the proposal head and $1e^{-5}$ for the backbone. We set the weight decay as $1e^{-4}$. MonTrack-T and MonTrack-B were trained with 57K and 81K steps, respectively. We empirically set the value of α as 1.0, and the masking ratios of RGB and event data δ_i and δ_e as 0.4 and 0.3, respectively.

We refer readers to the *Supplementary Material* for the

¹We refer readers to the *Supplementary Material* for more details

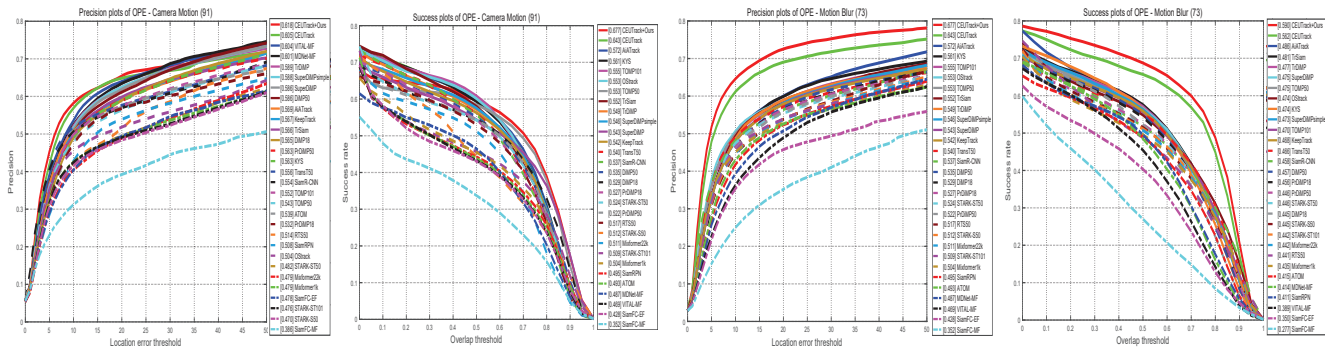


Figure 5. Visualization of the precision and success plots of sequences with different attributes in *COESOT* dataset. The set of camera motion and motion blur contains 91 and 73 sequences, respectively. We also refer readers to the *Supplementary Material* for comprehensive evaluations of all attributes [Zoom in to see details](#).

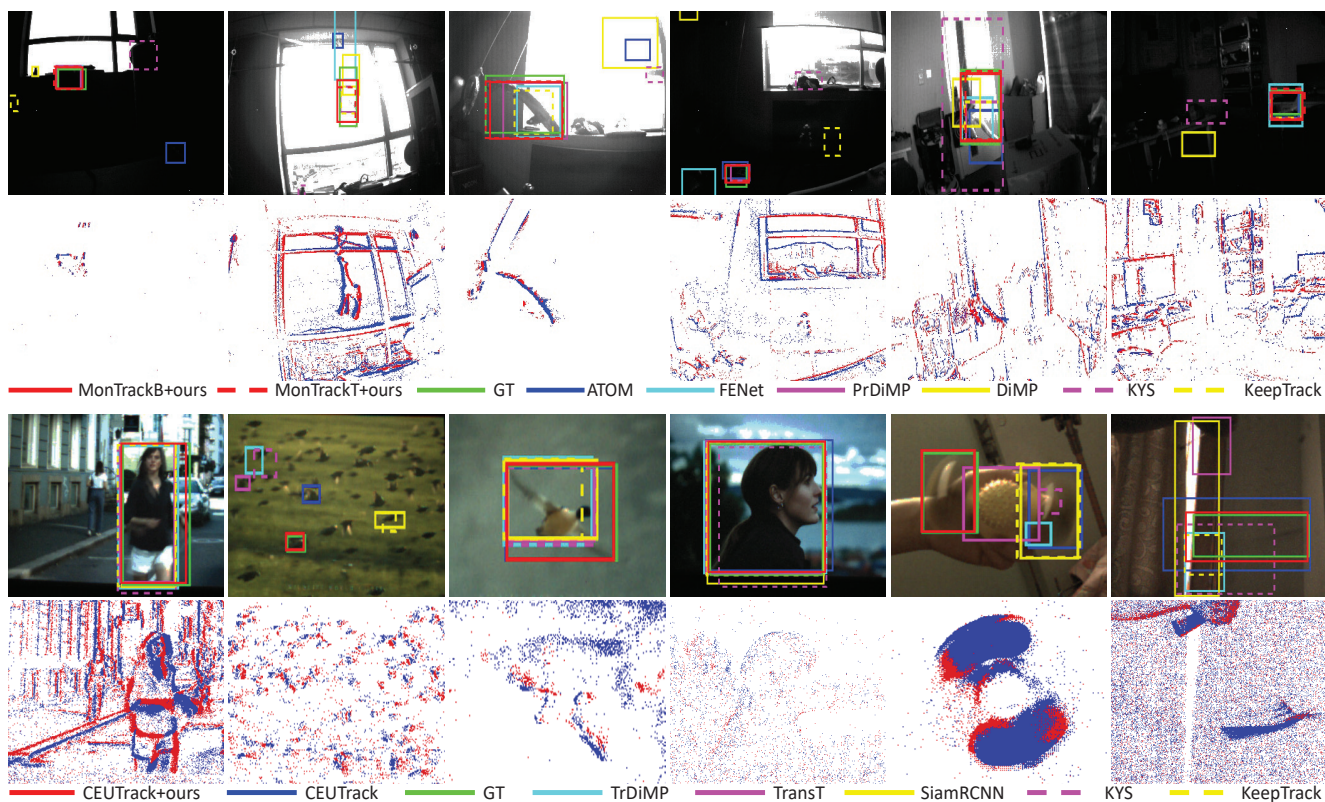


Figure 6. Visual comparisons of the tracking performance of different methods on the **(Upper)** FE108 and **(Bottom)** CEOSOT datasets. We also refer readers to the *video demo* for more visual comparisons.

by 1.1%. After being augmented with our techniques, i.e., CEUTrack+Ours, the values of SR and PR are further improved by 1.2% and 1.4%, respectively, and its NPR achieves higher than 70%, convincingly validating the effectiveness of the proposed techniques. In addition, we also provide the success and precision plots of different attributes in Fig. 6, where it can be seen that the proposed augmentations can yield general improvements instead of only strengthening certain circumstances. For example, the

proposed augmentations achieve 3.4 % precision and 2.8 % success improvements under the blurring attribute. Especially, CEUTrack+Ours maintains the best performance under the camera motion attribute, while the baseline CEUTrack drops to the 7th.

We also refer readers to the *Supplementary Material* for the comparisons of the network size and inference time.

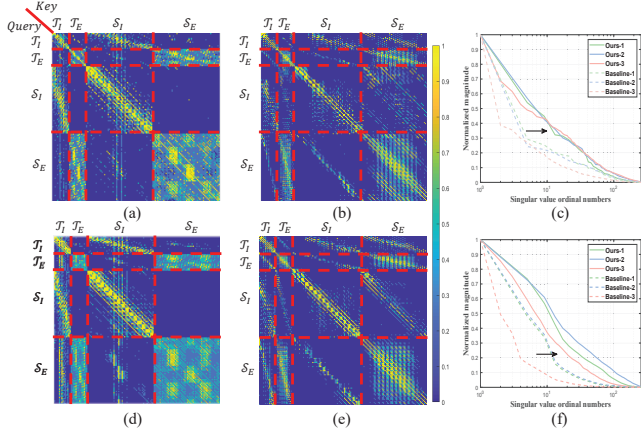


Figure 7. Visualization of the internal attention matrix, where we averaged six samples. (a) and (d) (resp. (b) and (e)) show the attention matrices of 5^{th} and 6^{th} attention layers of the baseline model CEUTrack (resp. the augmented model CEUTrack+Ours), respectively. (c) (resp. (f)) depicts the singular values of block matrices in (a) and (b) (resp. (d) and (e)) with “*-1”, “*-2”, and “*-3” indicating the block matrices of M_{T_I, T_E} , M_{T_I, S_E} , and M_{S_I, S_E} , respectively. We normalize both dimensions of curves in (c) and (f) into an identical range for better visualization. Note that we arranged these matrices with the same manner as that in Fig. 3(b); hence, the summation of each row equals 1. See the supplementary material for more visualizations of the orthogonal property.

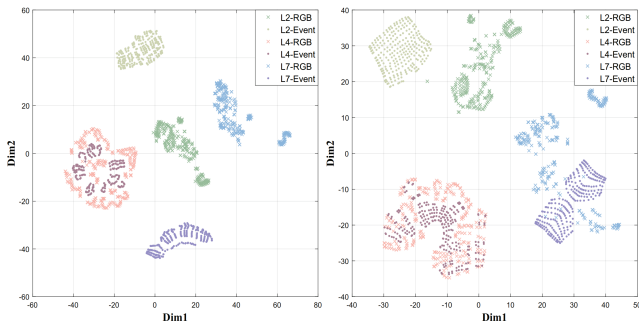


Figure 8. t-SNE visualization [48] of the **query vector** in self-attention with the 2^{nd} , 4^{th} , and 7^{th} layers. Visualization of the tokens from (Left) the baseline model CEUTrack and (Right) CEUTrack augmented with the proposed techniques during training.

4.2. Ablation Study

Visualizations. Fig. 7 visualizes the internal attention matrix of CEUTrack. The values of each row of the matrix are utilized to weight-sum tokens in that row and project to a corresponding token. Due to the absence of values in the blocks M_{S_I, S_E} , M_{S_I, T_E} , M_{T_I, T_E} , M_{T_I, S_E} in Figs. 7 (a) and (d), there is scarce information projected from the event domain to the RGB domain. The reason may be that the ViT was pre-trained on ImageNet composed of RGB data, making it preferable to process RGB data. When used as the backbone for constructing RGB-event object tracking,

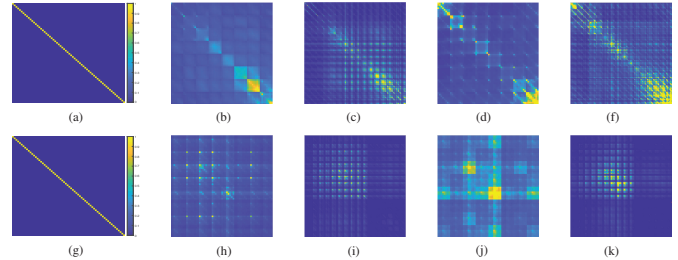


Figure 9. Visual comparisons of matrix orthogonality, where we visualize the matrix $A^T \cdot A$ of a given attention matrix $A \in \mathbb{R}^{n \times m}$ from CEUTracker on the COESOT dataset. (a) and (g) denote the fully-regularized orthogonal matrices, i.e., the optimization target. (b), (c), (d) and (f) (reps. (h), (i), (j) and (k)) indicate the matrices with (reps. without) the proposed augmentation scheme from M_{T_E, S_I} of 5^{th} , M_{S_E, S_I} of 5^{th} , M_{T_E, S_I} of 6^{th} , and M_{S_E, S_I} of 6^{th} layer, respectively.




Figure 10. Visual comparisons of matrix orthogonality, where we visualize the matrix $A^T \cdot A$ of a given attention matrix $A \in \mathbb{R}^{n \times m}$ from MonTrack-B on the FE108 dataset. (a) and (g) denote the fully-regularized orthogonal matrices, i.e., the optimization target. (b), (c), (d) and (f) (reps. (h), (i), (j) and (k)) indicate the matrices with (reps. without) the proposed augmentation scheme from $M_{T, S}$ of 3^{th} layer 4^{th} head, $M_{T, S}$ of 3^{th} layer 5^{th} head, $M_{T, S}$ of 3^{th} layer 6^{th} head, and $M_{T, S}$ of 4^{th} layer 8^{th} head, respectively.

the pre-trained filters attempt to project event information onto RGB tokens to complete the labor-intensive tasks of information fusion and processing, instead of the inverse projection. After being augmented with our techniques during training, the cross-modal interaction is noticeably enhanced, i.e., the matrix blocks, which are zeros in Figs. 7 (a) and (d), exhibit attention values, as demonstrated in Figs. 7 (b) and (e). Besides, we also visualized the singular values of matrix blocks related to the cross-modal interaction in Figs. 7 (c) and (f), which substantially validates they have been pushed far away from a low-rank matrix after applying the proposed techniques. We refer readers to the *Supplementary Material* for more results. Finally, Fig. 8 shows the queries of the 2^{nd} , 4^{th} , and 7^{th} self-attention layers where it can be seen that the proposed augmentations narrow the distribution gaps between event and RGB tokens, especially for the 4^{th} layer.

Additional visualization of attention matrices. To further directly visualize the variations of matrix orthogonality after applying the proposed augmentation, we visualized the

Table 3. Quantitative comparison of the model efficiency on the COESOT dataset.

Methods	RTS50	PrDiMP50	KYS	DiMP50	KeepTrack	TrSiam	AiATrack	OSTrack
FPS	30	30	20	43	18	35	38	105
Methods	ToMP101	TrDiMP	TransT	SuperDiMP	SiamR-CNN	CEUTrack	CEUTrack+Ours	
FPS	20	26	50	-	5	75	75	

matrix $M^T \times M$ in Figs. 9 and 10, convincingly demonstrating the effectiveness of our regularization for enhancing a matrix’s row-/column-wise orthogonality.

Masking vs. High-rank. We conducted throughout experiments to better understand the relationship and function of the proposed two augmentation techniques. From Table 4, it can be seen that when the two techniques were simultaneously applied, the improvement is much more significant than that of only applying the masking scheme. The improvement is slight when only the high-rank regularization was applied. These observations validate our claim that the two techniques are complementary.

Effect of the mask size. We experimentally validated the effect of different mask sizes on performance. As shown in Table 5, the benefits may be nullified under extremely large or tiny masks. The possible reason is that the network experiences the small masks as noise. While if the mask is too broad, the object may only appear in one modal, which may be detrimental to cross-modal learning.

Model Efficiency. We list the speed of different methods for comparison. Note that the proposed augmentation techniques were applied to baseline models only during the training phase. Consequently, no extra computational burden is imposed on the baseline models during testing. For the training phase, the CEUTrack operates at a rate of 30.6 samples per second. Moreover, it executes at 28.8 samples per second after plugging the presented regularization term.

4.3. Discussion

In view of the impressive performance of the proposed plug-and-play training augmentations, it is worth further exploring their potential in other cross-modal scenarios, such as RGB-3D point clouds, or even vision-natural language. In addition, as demonstrated in Fig. 7, the proposed orthogonal high-rank regularization indeed facilitates the interactions between cross-modal tokens, and thus, it would be promising to further develop task-specific regularization terms for other visual Transformers-based works.

5. Conclusion

In this paper, we introduced plug-and-play training augmentations for Transformer-based RGB-event object tracking. Our augmentations consist of two complementary techniques—cross-modal mask modeling and orthogonal high-rank regularization with the same objective of enhancing the cross-modal interaction of a ViT pre-trained only

Table 4. Results of the ablative study on the effect of the proposed two training augmentation techniques.

Baseline	Masking	High-rank	RSR	OP _{0.50}	OP _{0.75}	RPR
MonTrack-B	×	×	64.3	84.7	35.2	93.2
	✓	×	67.6	88.6	41.8	96.6
	×	✓	65.4	85.9	37.8	94.6
	✓	✓	68.5	89.4	45.4	96.2
Baseline	Masking	High-rank	SR	PR	NPR	BOC
CEUTrack	×	×	62.0	70.5	69.0	20.77
	✓	×	62.4	71.0	69.3	21.13
	×	✓	61.7	70.2	68.6	20.90
	✓	✓	63.2	71.9	70.2	21.58

Table 5. Results of the ablative study on the mask size.

Methods	Mask size	RSR	OP _{0.50}	OP _{0.75}	RPR
MonTrack-T	1/2	65.0	84.2	37.9	94.3
	1/4	64.8	84.6	34.0	95.8
	1/8	66.3	86.4	40.0	95.3
	1/16	64.2	84.5	35.5	93.6

with RGB data. Our extensive experiments demonstrate the effectiveness of our training augmentations, as state-of-the-art methods achieve significant improvement in tracking performance after augmentation.

While current Transformers can be scaled up to enormous sizes, relying solely on final objectives to guide the model learning process may be insufficient. We hope our perspectives, findings and analysis will inspire further research into the *internal mechanisms* of Transformer-based cross-modal fusion tasks.

References

- [1] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *Proc. of ECCV*, pages 850–865, 2016. 2
- [2] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *Proc. of the IEEE/CVF ICCV*, pages 6182–6191, 2019. 2, 6
- [3] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Know your surroundings: Exploiting scene information for object tracking. In *Proc. of ECCV*, pages 205–221. Springer, 2020. 2, 6
- [4] Samuel Bryner, Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. Event-based, direct camera tracking from

- a photometric 3d map using nonlinear optimization. In *Proc. of the IEEE ICRA*, pages 325–331, 2019. 2
- [5] Yujeong Chae, Lin Wang, and Kuk-Jin Yoon. Siamevent: Event-based object tracking via edge-aware similarity learning with siamese networks. *arXiv preprint*, 2021. 3
- [6] Boyu Chen, Peixia Li, Lei Bai, Lei Qiao, Qihong Shen, Bo Li, Weihao Gan, Wei Wu, and Wanli Ouyang. Backbone is all your need: a simplified architecture for visual object tracking. In *Proc. of ECCV*, pages 375–392. Springer, 2022. 2
- [7] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *Proc. of IEEE/CVF CVPR*, pages 8126–8135, 2021. 1, 3, 5, 6
- [8] Martin Danelljan and Goutam Bhat. Pytracking: Visual tracking library based on pytorch, 2019. 6
- [9] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Atom: Accurate tracking by overlap maximization. In *Proc. of the IEEE/CVF CVPR*, pages 4660–4669, 2019. 6
- [10] Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *Proc. of the IEEE/CVF CVPR*, pages 7183–7192, 2020. 6
- [11] Pierre de Tournemire, Davide Nitti, Etienne Perot, Davide Migliore, and Amos Sironi. A large scale event-based detection dataset for automotive. *arXiv preprint*, 2020. 1
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proc. of the IEEE/CVF CVPR*, pages 248–255. Ieee, 2009. 3
- [13] Xingping Dong, Jianbing Shen, Ling Shao, and Fatih Porikli. Clnet: A compact latent network for fast adjusting siamese trackers. In *Proc. of ECCV*, pages 378–395. Springer, 2020. 6
- [14] Alexey Dosovitskiy, Lucas Beyer, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proc. of ICLR*. 2, 4
- [15] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, et al. Event-based vision: A survey. *IEEE TPAMI*, 44(1):154–180, 2020. 1
- [16] Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In *Proc. of CVPR*, pages 3867–3876, 2018. 1
- [17] Shenyan Gao, Chunluan Zhou, Chao Ma, Xinggang Wang, and Junsong Yuan. Aiatrack: Attention in attention for transformer visual tracking. In *Proc. of ECCV*, pages 146–164. Springer, 2022. 6
- [18] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. Asynchronous, photometric feature tracking using events and frames. In *Proc. of ECCV*, pages 750–765, 2018. 1
- [19] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. Ekl: Asynchronous photometric feature tracking using events and frames. *IJCV*, 128(3):601–618, 2020. 1
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proc. of the IEEE/CVF CVPR*, pages 16000–16009, 2022. 4
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. of IEEE/CVF CVPR*, pages 770–778, 2016. 2
- [22] Maximilian Jaritz, Tuan-Hung Vu, Raoul de Charette, Emilie Wirbel, and Patrick Pérez. xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation. In *Proc. of the IEEE/CVF CVPR*, pages 12605–12614, 2020. 3
- [23] Sajid Javed, Martin Danelljan, Fahad Shahbaz Khan, Muhammad Haris Khan, Michael Felsberg, and Jiri Matas. Visual object tracking with discriminative filters and siamese networks: a survey and outlook. *IEEE TPAMI*, 2022. 2
- [24] Ilchae Jung, Jeany Son, Mooyeol Baek, and Bohyung Han. Real-time mdnet. In *Proc. of ECCV*, pages 83–98, 2018. 2
- [25] Simon Klenk, David Bonello, Lukas Koestler, and Daniel Cremers. Masked event modeling: Self-supervised pretraining for event cameras. *arXiv preprint*, 2022. 4
- [26] Beat Kueng, Elias Mueggler, Guillermo Gallego, and Davide Scaramuzza. Low-latency visual odometry using event-based feature tracks. In *Proc. of the IEEE/RSJ IROS*, pages 16–23. IEEE, 2016. 1
- [27] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *Proc. of the IEEE/CVF CVPR*, pages 4282–4291, 2019. 2
- [28] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8971–8980, 2018. 2
- [29] Hongmin Li and Luping Shi. Robust event-based object tracking combining correlation filter and cnn representation. *Frontiers in neurorobotics*, 13:82, 2019. 3
- [30] Jianing Li, Jia Li, Lin Zhu, Xijie Xiang, Tiejun Huang, and Yonghong Tian. Asynchronous spatio-temporal memory network for continuous event-based object detection. *IEEE TIP*, 31:2975–2987, 2022. 1
- [31] Liting Lin, Heng Fan, Zhipeng Zhang, Yong Xu, and Haibin Ling. Swintrack: A simple and strong baseline for transformer tracking. In *Proc. NeurIPS*. 1, 2
- [32] Huayao Liu, Jiaming Zhang, Kailun Yang, Xinxin Hu, and Rainer Stiefelhagen. Cmx: Cross-modal fusion for rgb-x semantic segmentation with transformers. *arXiv preprint arXiv:2203.04838*, 2022. 3
- [33] Ze Liu, Han Hu, et al. Swin transformer v2: Scaling up capacity and resolution. In *Proc. of the IEEE/CVF CVPR*, pages 12009–12019, 2022. 5
- [34] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proc. of the IEEE/CVF ICCV*, pages 10012–10022, 2021. 2, 5
- [35] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proc. of ICLR*. 5
- [36] Christoph Mayer, Martin Danelljan, Goutam Bhat, Matthieu Paul, Danda Pani Paudel, Fisher Yu, and Luc Van Gool.

- Transforming model prediction for tracking. In *Proc. of the IEEE/CVF CVPR*, pages 8731–8740, 2022. 6
- [37] Christoph Mayer, Martin Danelljan, Danda Pani Paudel, and Luc Van Gool. Learning target candidate association to keep track of what not to track. In *Proc. of IEEE/CVF ICCV*, pages 13444–13454, 2021. 6
- [38] Anton Mitrokhin, Cornelia Fermüller, Chethan Parameshwara, and Yiannis Aloimonos. Event-based moving object detection and tracking. In *Proc. of IEEE/RSJ IROS*, pages 1–9. IEEE, 2018. 1, 2
- [39] Anindya Mondal, Jhony H Giraldo, Thierry Bouwmans, Ananda S Chowdhury, et al. Moving object detection for event-based vision using graph spectral clustering. In *Proc. of the IEEE/CVF ICCV*, pages 876–884, 2021. 1
- [40] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *Proc. of the IEEE/CVF CVPR*, pages 4293–4302, 2016. 2
- [41] Matthieu Paul, Martin Danelljan, Christoph Mayer, and Luc Van Gool. Robust visual tracking by segmentation. In *Proc. of ECCV*, pages 571–588. Springer, 2022. 6
- [42] Etienne Perot, Pierre de Tournemire, Davide Nitti, Jonathan Masci, and Amos Sironi. Learning to detect objects with a 1 megapixel event camera. In *Proc. of NeurIPS*, volume 33, pages 16639–16652, 2020. 1
- [43] Dhanesh Ramachandram and Graham W Taylor. Deep multimodal learning: A survey on recent advances and trends. *IEEE signal processing magazine*, 34(6):96–108, 2017. 3
- [44] Bharath Ramesh, S. Zhang, Zhi Wei Lee, Zhi Gao, Garrick Orchard, and Cheng Xiang. Long-term object tracking with a moving event camera. In *Proc. of BMVC*, 2018. 2
- [45] Bharath Ramesh, Shihao Zhang, Hong Yang, Andres Ussa, Matthew Ong, Garrick Orchard, and Cheng Xiang. e-tld: Event-based framework for dynamic object tracking. *IEEE TCSVT*, 31(10):3996–4006, 2020. 1
- [46] S Ren, K He, R Girshick, and J Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE TPAMI*, 39(6):1137–1149, 2016. 2
- [47] Chuanming Tang, Xiao Wang, Ju Huang, Bo Jiang, Lin Zhu, Jianlin Zhang, Yaowei Wang, and Yonghong Tian. Revisiting color-event based tracking: A unified network, dataset, and metric. *arXiv preprint arXiv:2211.11010*, 2022. 4, 5, 6
- [48] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 9(11), 2008. 8
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Proc. of NeurIPS*, 30, 2017. 4
- [50] Paul Voigtlaender, Jonathon Luiten, Philip HS Torr, and Bastian Leibe. Siam r-cnn: Visual tracking by re-detection. In *Proc. of the IEEE/CVF CVPR*, pages 6578–6588, 2020. 6
- [51] Chunwei Wang, Chao Ma, Ming Zhu, and Xiaokang Yang. Pointaugmenting: Cross-modal augmentation for 3d object detection. In *Proc. of the IEEE/CVF CVPR*, pages 11794–11803, 2021. 3
- [52] Chaoqun Wang, Chunyan Xu, Zhen Cui, Ling Zhou, Tong Zhang, Xiaoya Zhang, and Jian Yang. Cross-modal pattern-propagation for rgb-t tracking. In *Proc. of the IEEE/CVF CVPR*, pages 7064–7073, 2020. 3
- [53] Guangting Wang, Chong Luo, Xiaoyan Sun, Zhiwei Xiong, and Wenjun Zeng. Tracking by instance detection: A meta-learning approach. In *Proc. of the IEEE/CVF CVPR*, pages 6288–6297, 2020. 2
- [54] Ning Wang, Wengang Zhou, Jie Wang, and Houqiang Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *Proc. of the IEEE/CVF CVPR*, pages 1571–1580, 2021. 1, 3, 6
- [55] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proc. of the IEEE/CVF CVPR*, pages 9653–9663, 2022. 4
- [56] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *Proc. of the IEEE/CVF ICCV*, pages 10448–10457, 2021. 2
- [57] Botao Ye, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Joint feature learning and relation modeling for tracking: A one-stream framework. In *Proc. of ECCV*, pages 341–357. Springer, 2022. 2, 6
- [58] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-modal self-attention network for referring image segmentation. In *Proc. CVPR*, pages 10502–10511, 2019. 3
- [59] Jiqing Zhang, Bo Dong, Haiwei Zhang, Jianchuan Ding, Felix Heide, Baocai Yin, and Xin Yang. Spiking transformers for event-based single object tracking. In *Proc. of the IEEE/CVF CVPR*, pages 8801–8810, 2022. 1
- [60] Jiqing Zhang, Xin Yang, Yingkai Fu, Xiaopeng Wei, Baocai Yin, and Bo Dong. Object tracking by jointly exploiting frame and event domain. In *Proc. of the IEEE/CVF ICCV*, pages 13043–13052, 2021. 1, 6
- [61] Yifan Zhang, Qijian Zhang, Junhui Hou, Yixuan Yuan, and Guoliang Xing. Bidirectional propagation for cross-modal 3d object detection. *arXiv preprint arXiv:2301.09077*, 2023. 3
- [62] Changqing Zhou, Zhipeng Luo, Yueru Luo, Tianrui Liu, Liang Pan, Zhongang Cai, Haiyu Zhao, and Shijian Lu. Ptrr: Relational 3d point cloud object tracking with transformer. In *Proc. of the IEEE/CVF CVPR*, pages 8531–8540, 2022. 1
- [63] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Unsupervised event-based learning of optical flow, depth, and egomotion. In *Proc. of the IEEE/CVF CVPR*, pages 989–997, 2019. 1
- [64] Zhiyu Zhu, Junhui Hou, and Xianqiang Lyu. Learning graph-embedded key-event back-tracing for object tracking in event clouds. In *Proc. of NeurIPS*, 2022. 1, 3, 5, 6
- [65] Zhiyu Zhu, Hui Liu, Junhui Hou, Huanqiang Zeng, and Qingfu Zhang. Semantic-embedded unsupervised spectral reconstruction from single rgb images in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2279–2288, 2021. 3