

What You Have is What You Track: Adaptive and Robust Multimodal Tracking

Yuedong Tan^{1,2,3} Jiawei Shao^{1*} Eduard Zamfir² Ruanjun Li⁴ Zhaochong An⁵
 Chao Ma⁶ Danda Paudel³ Luc Van Gool³ Radu Timofte² Zongwei Wu^{2*}

¹ TeleAI, China Telecom ² Computer Vision Lab, CAIDAS & IFI, University of Wurzburg ³ INSAIT, Sofia University
⁴ ShanghaiTech University ⁵ University of Copenhagen ⁶ AI Institute, Shanghai Jiao Tong University

Abstract

Multimodal data is known to be helpful for visual tracking by improving robustness to appearance variations. However, sensor synchronization challenges often compromise data availability, particularly in video settings where shortages can be temporal. Despite its importance, this area remains underexplored. In this paper, we present the first comprehensive study on tracker performance with temporally incomplete multimodal data. Unsurprisingly, under such a circumstance, existing trackers exhibit significant performance degradation, as their rigid architectures lack the adaptability needed to effectively handle missing modalities. To address these limitations, we propose a flexible framework for robust multimodal tracking. We venture that a tracker should dynamically activate computational units based on missing data rates. This is achieved through a novel Heterogeneous Mixture-of-Experts fusion mechanism with adaptive complexity, coupled with a video-level masking strategy that ensures both temporal consistency and spatial completeness — critical for effective video tracking. Surprisingly, our model not only adapts to varying missing rates but also adjusts to scene complexity. Extensive experiments show that our model achieves SOTA performance across 9 benchmarks, excelling in both conventional complete and missing modality settings. The code and benchmark will be made publicly available at <https://github.com/supertyd/FlexTrack>.

1. Introduction

Visual object tracking aims to localize a target object in a video while handling dynamic appearance changes [3, 5, 15, 33, 39, 53]. To improve robustness against occlusions, illumination shifts, and environmental variations, multimodal tracking has emerged as a powerful solution [29, 50, 52]. Complementary information from multiple sensors (e.g., depth, thermal, and event) has shown the ef-

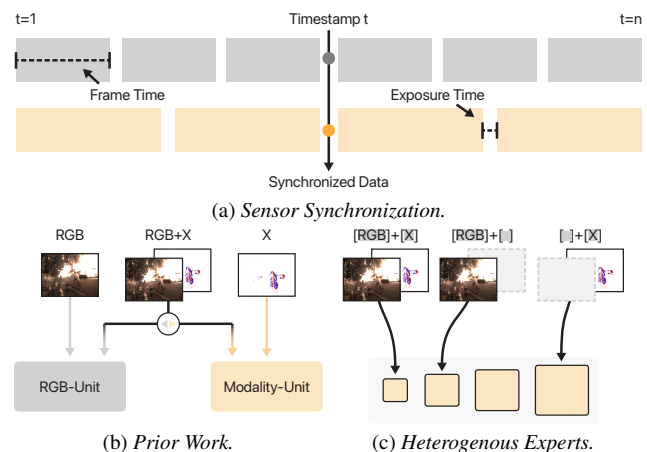


Figure 1. Synchronization has long been a challenge in multisensor fusion. As illustrated in (a), differences in frame timing and exposure duration make it difficult for a multisensor system to obtain perfectly synchronized pairs, often resulting in missing data. To address this issue, existing trackers [36], as shown in (b), handle each missing-data scenario individually, employing tailored and individual strategies for different cases. However, this approach does not account for the complexity variations introduced by different types of missing data. In contrast, in (c), we take a more holistic approach, unifying both missing and full-modality scenarios within a single framework. Our method leverages Heterogeneous Experts with adaptive complexity, providing a personalized yet structurally unified solution for robust tracking.

fectiveness of improving unimodal approaches under ideal conditions.

However, real-world multimodal data is rarely perfect [37, 38, 66]. Sensors often suffer from synchronization failures and temporary dropouts. As illustrated in Fig. 1, the inherent differences in exposure time and frame rates make it extremely challenging to achieve perfect synchronization across multiple sensors. Existing tracking datasets [29, 31, 60] simplify this issue by considering only paired modalities, such as RGB and depth. However, this setting remains limited, as it does not generalize to real-world scenarios where multiple modalities may be intermittently unavailable within a given temporal window during video tracking. Despite its practical importance, such a missing

*Corresponding

Table 1. Comparison of SOTA multimodal tracking models.

Method	Unified Arch.	Unified Param.	Missing Modal	Flexi-Modal Adaptive
IPT [36]	✗	✗	✓	✗
ViPT [67]	✓	✗	✗	✗
UnTrack [56]	✓	✓	✗	✗
Ours	✓	✓	✓	✓

scenario has not been well-adjusted.

In this paper, we conduct a comprehensive study on temporally missing modalities in multimodal tracking [6, 21, 67]. For each benchmark dataset, we introduce a missing-data variation that simulates real-world synchronization challenges by incorporating diverse missing patterns, excluding cases where all modalities are absent simultaneously. Benchmarking existing methods under this setting reveals that they struggle to maintain reliable performance, primarily due to their reliance on template-matching frameworks between a reference template and a search region [5, 18, 32, 33, 62, 67], which are not designed to handle missing modalities. A recent attempt [36] to address this issue introduces placeholder and tailored prompts for both complete and missing cases. However, its rigid design raises a fundamental question: *Should a tracker maintain the same computational complexity regardless of whether modalities are missing or available?*

Intuitively, we argue that an effective tracker should dynamically adapt, adjusting the computational resources it uses based on the available information and the severity of missing data. To achieve this, we propose **FlexTrack**, a unified framework designed to robustly handle multimodal tracking in both full and missing modality scenarios. Technically, we introduce a novel heterogeneous Mixture-of-Experts(HMoE) mechanism. This approach leverages varying levels of complexity and offers a unique advantage by routing and activating the most relevant experts based on the current data. Additionally, we propose a novel video-level masking strategy that enforces temporal consistency at the clip level while preserving contextual completeness, thus facilitating more robust object tracking.

Our tracker stands out from SOTA counterparts, as demonstrated by the holistic comparison in Tab. 1, making it more adaptable to tracking with available data. Extensive evaluations on both conventional benchmarks with full modalities and our newly created missing-data variations validate its effectiveness. Notably, on full-modality benchmarks, we outperform the current SOTA by 2.6%, with the performance gap widening significantly under the missing-data setting.

2. Related Work

2.1. Multimodal Tracking

Modern object tracking has advanced significantly with the adoption of Siamese networks and transformer architectures [5, 16, 53, 62]. Despite these advancements, track-

ing remains constrained by the limitations of single-modal (e.g., RGB) inputs, which struggle with challenges like occlusion or low-light conditions [51]. Multimodal tracking [22, 29, 69], which integrates RGB data with additional modalities (e.g., depth, thermal, event), has proven effective, especially when temporal cues are exploited [20, 21, 56, 67]. However, these methods assume consistent availability of modality, which is unrealistic in real-world scenarios [36]. Our work addresses this gap by introducing a video-level masking strategy that explicitly handles missing modalities while retaining video-level temporal coherence, and a HMoE mechanism to switch different sizes of experts to handle different levels of modality missing.

2.2. Missing Modalities

Missing modalities [37, 38, 66] is a widely researched the problem. Real-world multimodal systems often face the challenge of having to handle cases where certain data modalities may be missing or incomplete. Existing methods for addressing missing modalities in object tracking include prompt-based approaches and generative models such as GANs [17] or diffusion models [8, 24]. While effective, these methods often introduce additional computational overhead. Alternatively, some approaches employ distillation techniques [1, 48], utilizing full-modal models to distill models with missing modalities. However, this can result in reduced model flexibility. Another common strategy involves actively masking modalities during training to enhance the model’s robustness to each modality [54]. This technique typically involves randomly masking different modalities to improve the model’s resilience. However, existing models primarily focus on single-instance prediction tasks such as classification and detection. In contrast, our model explores a novel video-level masking strategy for object tracking, effectively leveraging temporal information to mitigate the impact of missing modalities.

2.3. Mixture-of-Experts

Mixture-of-Experts(MoE) models utilize multiple specialized expert networks to adaptively select and process inputs, thereby enhancing performance across various tasks. Notably, large-scale MoE implementations such as DeepSeek-MoE [9] and Llama-MoE [68] have demonstrated exceptional efficiency in performance. Beyond their prominence in large language models, MoE architectures have also been effectively applied in multimodal fusion, leveraging flexible expert switching to accommodate diverse scenarios. For instance, FuseMoE [19] and Flex-MoE [63] employ MoE to integrate different modalities and address issues arising from missing modalities. However, they all have the uniform design for the expert size.

In contrast, our approach introduces a novel module by combining temporal information with experts of varying

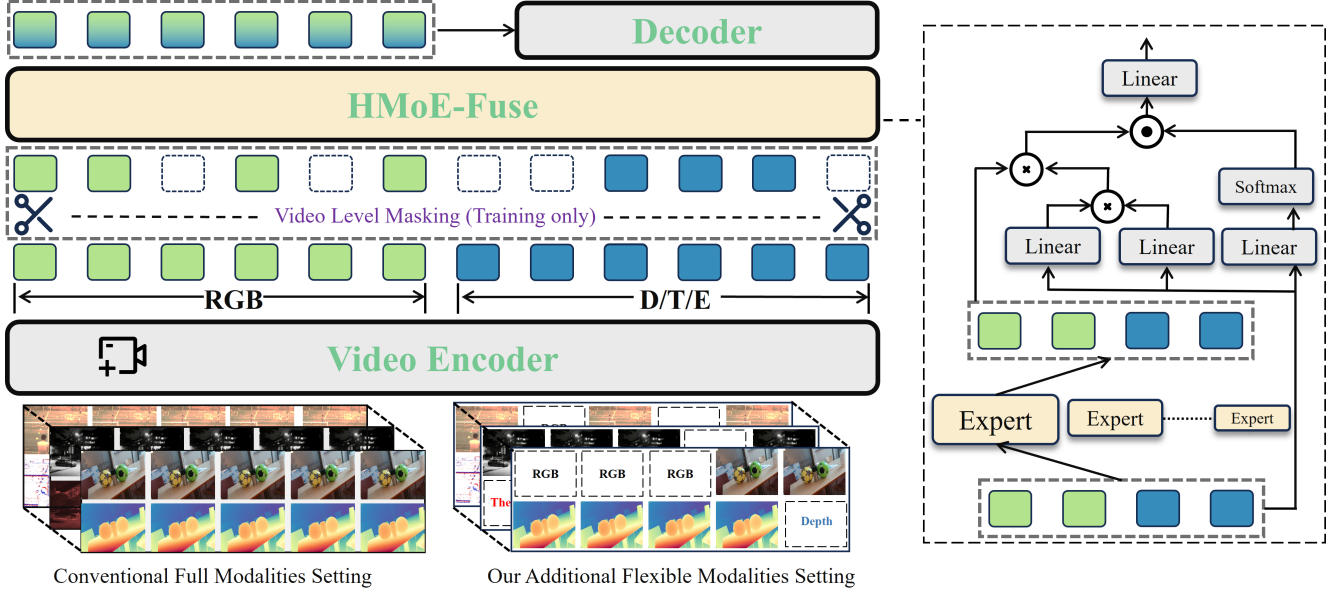


Figure 2. **Our framework:** Conventional trackers operate under the strong assumption that all frames are always available, disregarding practical scenarios with missing modalities. To address this limitation, we introduce heterogeneous experts, which dynamically switch between different experts to adapt *test-time complexity*. To further strengthen the robust learning, we employ a video-level masking strategy during training, encouraging the model to capture temporal continuity despite missing data.

sizes to address the challenges of adaptive fusion. This methodology enables more precise and context-aware processing of multimodal data, thereby enhancing the model’s robustness and flexibility in handling diverse scenarios.

3. Method

3.1. Overview

In this paper, we propose a unified framework for object tracking that can handle flexible combinations of modalities. As illustrated in Fig. 2, our framework comprises two primary components: a multimodal fusion module and a video-level masking strategy which is only used in the training stage. Our tracker accepts video clips and search regions from RGB and auxiliary modalities (depth, thermal, and event modalities, collectively referred to as \mathbf{X}) as input.

To address the challenge of varying modality completeness in tracking, we propose the Heterogeneous Mixture-of-Experts Fusion (HMoE-Fuse) module. Unlike conventional approaches that statically allocate computational resources, we dynamically adjust the model’s test-time complexity based on the availability of multimodal inputs.

Additionally, to enhance temporal continuity learning under partial modality-missing conditions, we propose a multimodal video-level masking mechanism during training. In each time step, we keep at least one frame available to focus on learning spatio-temporal consistency.

3.2. Heterogeneous Mixture-of-Experts Fusion

As illustrated in Fig. 2, we first feed both the RGB and \mathbf{X} frames into the video encoder, which produces several token representations: RGB search region tokens, denoted as T_{rgb}^s ; video clip tokens ranging from T_{rgb}^{c1} to T_{rgb}^{ci} ; and, for additional \mathbf{X} , we obtain T_x^s, T_x^{c1} to T_x^{ci} , where $i \in \{1, \dots, N\}$.

First, we concatenate all tokens to form the sequence token T_v . The core component of the HMoE-Fuse is a heterogeneous MoE, which differs from typical MoEs where each expert is constructed with a distinct architecture. Specifically, each expert in our model consists of a linear layer, and its hidden state dimension is set as 2^d , where $d \in \{2, \dots, D-1\}$.

Contrary to prior MoEs [9, 19], we route the entire video clips to experts instead of individual tokens or images. First, a gating function G produces a weight distribution, where higher values indicate better input-expert alignment. The output tokens T_y^1 of the HMoE layer is the weighted sum of outputs from the top- K activated experts:

$$g_n = \begin{cases} \text{Softmax}(G(T_v))_n, & \text{if } G(T_v)_n \in \text{Top-K}(G(T_v)) \\ 0, & \text{otherwise} \end{cases}$$

$$T_y^1 = \sum_{n=1}^M g_n E_n(T_v). \quad (1)$$

where g_n represents the gating values derived from the function $G(T_v)$, and $E_n(T_v)$ is the output of the n -th expert. Besides, the gating function ensures that only the top K ex-

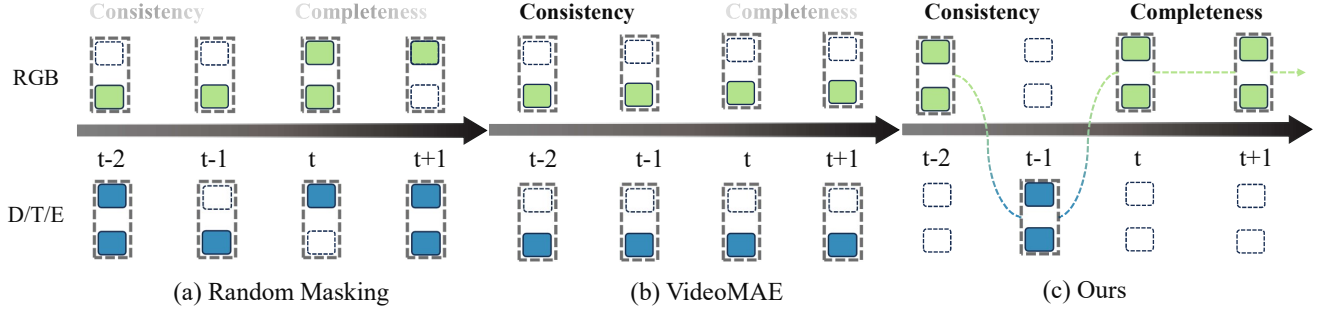


Figure 3. Comparison of different masking strategies for multimodal settings. (a) Conventional random or MAE-like masking [55] can be easily extended to multimodal settings but disrupts both spatial completeness and temporal continuity in video frames. (b) Extending VideoMAE [45] to multimodal settings enhances temporal consistency by preserving token continuity within each modality. However, it applies the same spatial randomness across all frames, resulting in persistent spatial information loss. (c) We argue that both spatial completeness and temporal consistency are crucial for object tracking: spatial completeness provides a detailed understanding of appearance, while temporal consistency captures dynamic changes. To achieve this balance, we propose a structured, video-level masking strategy that ensures the full spatial representation remains consistently retrievable in available data across different timestamps.

perts are activated, here we set $K=2$ in our method. With consistent video frame input, our HMoE selects the most appropriate experts to deal with missing modality cases. Simultaneously, we process the tokens T_v through two linear layers, yielding an intermediate transformation:

$$T_y^2 = T_v W_1 (T_v W_2)^T, \quad (2)$$

where W_1 and W_2 are learned weight matrices that perform linear attention on the tokens.

Next, we perform a further transformation on the tokens by combining the expert outputs and the intermediate result:

$$T_y^3 = T_y^1 (T_y^2)^T. \quad (3)$$

This operation integrates the expert outputs with the transformed tokens T_y^2 . Finally, we apply a softmax operation and a linear transformation to further refine the output:

$$T_y^4 = \text{Softmax}(T_v W_3) T_y^3, \quad (4)$$

and

$$T_y^5 = T_y^4 W_4, \quad (5)$$

where W_3 and W_4 are additional learnable weight matrices. The softmax operation is applied to $T_v W_3$ to scale the transformed tokens, and W_4 further adjusts the resulting values.

3.3. Multimodal Video-Level Masking Strategy

During inference, we find that our multimodal tracker fail to handle the missing modality for its substantial dependence on each frame in the video clip. (For a correspondence experiment, please refer to Tab. 6.)

While VideoMAE [45]’s tube masking strategy demonstrates effectiveness in multiple video tasks through video cube masking, its design fundamentally assumes access to subsequent frames for masked content inference. However, the assumption is invalid for real-time tracking as future frames are inaccessible. Specifically, models cannot

Algorithm 1: Video-level masking

Input: RGB and X search region tokens T_{rgb}^s, T_x^s ,
 RGB and X video clip tokens T_{rgb}^{ci}, T_x^{ci} ,
 $i \in \{1, \dots, N\}$, batch size B , video clip
 mask ratio α .

Output: Updated tokens with applied video-level
 masking.

```

1 for  $b = 1$  to  $B$  do
2    $M_s \leftarrow \text{random}([1, 1], [1, 1], [1, 1], [1, 0], [0, 1]);$ 
3    $T_x^s \leftarrow M_s[0] \times T_x^s;$ 
4    $T_{rgb}^s \leftarrow M_s[1] \times T_{rgb}^s;$ 
5    $p \sim \mathcal{U}(0, 1);$  // Randomly select  $X$ 
   from the interval  $[0, 1];$ 
6   if  $p < \alpha$  then
7     for  $i = 1$  to  $N$  do
8        $M_{vi} \leftarrow \text{random}([1, 1], [1, 0], [0, 1]);$ 
9        $T_x^{ci} \leftarrow M_{vi}[0] \times T_x^{ci};$ 
10       $T_{rgb}^{ci} \leftarrow M_{vi}[1] \times T_{rgb}^{ci};$ 
11     end
12   end
13 end
```

utilize future temporal information to compensate for current modality absence, which poses a critical challenge in streaming tracking scenarios. We identify a unique advantage in multimodal tracking: temporal continuity persists even under spatial dropout in a modality. While individual frames may suffer from partial modality absence (e.g., missing depth data in frame t), adjacent frames ($t - \delta$ to $t + \delta$) inherently preserve temporal coherence through at least one persistent modality information. As shown in Fig. 3, we exploit this through a multimodal video-level masking strategy to enforce temporal modeling.

We illustrate the proposed video-level masking strategy in Algorithm 1. During training, we first apply search re-

gion masking by randomly selecting one of five predefined modality patterns to partially mask the RGB/X tokens in the search region. After applying the search region mask, a probability p is sampled from a uniform distribution over the interval $[0, 1]$. If the value of p falls below the predetermined threshold of α , this signifies the conditions necessary for implementing video-level masking. Thus, the algorithm proceeds to generate individual clip-level masks $M_{vi}, i \in 1, \dots, N$ for each of the video clips.

These masks are applied to the corresponding tokens of both modalities, further refining the focus of the model on the most pertinent information.

3.4. Head and Loss Function

To align with current multimodal trackers [20, 21, 56, 67], we utilize both a classification head and a regression head to localize the object center and approximate its height and width. Consequently, we adopt the same loss functions as these trackers, specifically the classification loss \mathcal{L}_{cls} , the L1 loss \mathcal{L}_{l1} , and the GIoU loss \mathcal{L}_{GIoU} . Beyond the standard tracking task losses, we integrate the importance loss $\mathcal{L}_{important}$ and balance loss $\mathcal{L}_{balance}$ from the Switch Transformer [14] to ensure load balancing among experts and prevent imbalance within a single expert. The auxiliary loss \mathcal{L}_{aux} is defined as follows:

$$\mathcal{L}_{aux} = \mathcal{L}_{balance} + \mathcal{L}_{important}. \quad (6)$$

The overall loss function is formulated as:

$$\mathcal{L}_{all} = \lambda_1 \mathcal{L}_{aux} + \lambda_2 \mathcal{L}_{cls} + \lambda_3 \mathcal{L}_{l1} + \lambda_4 \mathcal{L}_{GIoU}, \quad (7)$$

where \mathcal{L}_{all} represents the total loss used in the training process of our method. $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ is the hyperparameter with default values $\lambda_1 = 1, \lambda_2 = 5, \lambda_3 = 2$, and $\lambda_4 = 1$.

4. Experiment

4.1. Implementation Detail

Consistent with recent trackers [20–22], our tracker is initialized with a pretrained encoder [25, 44] trained on large-scale datasets [13, 23, 34, 40]. The model is trained on 8 4090 GPUs with a batch size of 32 using the AdamW optimizer, with a base learning rate of $3e-4$. Training consists of 30 epochs, each containing 600,000 image pairs. The learning rate is reduced by a factor of 10 after 22 epochs. For RGB-Thermal, we use LasHeR [31] for training, VisEvent [51] for RGB-Event, and DepthTrack [60] for RGB-Depth. To create a unified framework, we jointly train these three datasets in a single training process. During inference, we disable masking operations to fully leverage all available modalities while retaining the robustness learned through training phase masking simulations.

Table 2. SOTA comparisons on RGB-Event tracking dataset.

Method	VisEvent [51]	
	P	AUC
FlexTrack	81.4	64.1
STTrack [22]	78.6	61.9
SUTrack [6]	79.9	62.7
OneTracker [20]	76.7	60.8
SDSTrack [21]	76.7	59.7
UnTrack [56]	75.5	58.9
ViPT [67]	75.8	59.2
ProTrack [61]	63.2	47.1
MCITrack [25]	69.7	53.3
OSTrack [62]	69.5	53.4
SiamRCNN_E [46]	65.9	49.9
TransT_E [4]	65.0	47.4
LTMU_E [10]	65.5	45.9
PrDiMP_E [12]	64.4	45.3
VITAL_E [43]	64.9	41.5
MDNet_E [41]	66.1	42.6
ATOM_E [11]	60.8	41.2
SiamCar_E [59]	59.9	42.0
SiamBAN_E [7]	59.1	40.5
SiamMask_E [49]	56.2	36.9

4.2. Evaluations on complete dataset:

To demonstrate the performance of our tracker, we evaluated our tracker in 5 complete datasets. Under a strong assumption that all the frames are aligned well in temporal and there is not any missing modality caused by communication disruption and so on.

Evaluations on complete RGB-Event dataset: The VisEvent[51] test set contains 320 videos covering a variety of challenging attributes. Our trackers achieved new SOTA results, exceeding current SOTA trackers by 1.5% and 1.3% in P and area-under-the-curve (AUC), respectively.

Evaluations on complete RGB-Thermal dataset: For the RGB-Thermal tracking task, in Tab. 2, our FlexTrack models achieve new SOTA performance on both the LasHeR [31] and RGBT234 [29] datasets. Specifically, on the LasHeR benchmark, our FlexTrack achieves an AUC score of 62.0%, surpassing the previous best-performing model, STTrack [22], by 1.7%. On the RGBT234 benchmark, FlexTrack sets new SOTA score for MSR at 69.9%.

Evaluations on complete RGB-Depth dataset: In comprehensive RGB-Depth datasets, our model has demonstrated exceptional performance in DepthTrack [60]. Specifically, we have achieved an F-score that surpasses the previous best tracker, SUTrack [6], by 1.9%. Furthermore, despite training exclusively on the VisEvent [51], LasHeR [31], and VisEvent [51] datasets, our model exhibits strong out-of-distribution performance. In Tab. 4, experimental re-

Table 3. SOTA comparisons on RGB-Thermal tracking datasets.

Method	LasHeR [31]		RGBT234 [29]	
	P	AUC	MPR	MSR
FlexTrack	77.3	62.0	92.7	69.9
STTrack [22]	76.0	60.3	89.8	66.7
SUTrack [6]	74.5	59.5	92.2	69.5
OneTracker [20]	67.2	53.8	85.7	64.2
SDSTrack [21]	66.5	53.1	84.8	62.5
UnTrack [56]	64.6	51.3	84.2	62.5
ViPT [67]	65.1	52.5	83.5	61.7
ProTrack [61]	53.8	42.0	79.5	59.9
MCITrack [25]	64.2	50.2	79.1	60.4
OSTrack [62]	51.5	41.2	72.9	54.9
APFNet [57]	50.0	36.2	82.7	57.9
CMPP [47]	-	-	82.3	57.5
JMMAC [65]	-	-	79.0	57.3
CAT [30]	45.0	31.4	80.4	56.1
FANet [70]	44.1	30.9	78.7	55.3
MaCNet [64]	-	-	79.0	55.4

sults indicate that our model also achieves impressive results on the VOT-RGBD22 benchmark.

4.3. Evaluations on missing modality dataset

IPT [36] is the first method to address scenarios involving missing modalities. Initially, IPT was trained and tested exclusively on RGB-Thermal datasets. Subsequently, to better align with real-world applications, IPT developed a new dataset derived from the original one. This dataset includes instances of random missing, switched missing, and prolonged missing modalities. Building upon this dataset design, we expanded our evaluation to encompass multiple modalities, including DepthTrack [60], VisEvent [51], LasHeR [31], and RGBT234 [29].

RGB-Event missing modality datasets: As illustrated in Tab. 5, our tracker achieves a P of 72.8%, outperforming SUTrack [6] by 6.2%. Thanks to the video-level masking strategy, our model can capture more temporal information. This enables it to benefit from both the missing spatial information and the complete temporal information, even when the modalities are missing.

RGB-Thermal missing modality datasets: To evaluate the scenarios involving missing modalities on the RGB-Thermal benchmark, we select two datasets, namely LasHeR [31] and RGBT234 [29]. As shown in Tab. 5, IPT [36] is a tracker specifically designed for the RGB-Thermal missing modality. It requires additional prompt parameters for tracking and needs to switch among three branches to handle different inference scenarios, including RGB-only, Thermal-only, and RGB-Thermal. In contrast, our tracker does not introduce new parameters for miss-

Table 4. SOTA comparisons on RGB-Depth tracking datasets.

Method	DepthTrack [60]			VOT-RGBD22 [28]		
	F-score	Re	Pr	EAO	Acc.	Rob.
FlexTrack	67.0	66.9	67.1	78.0	83.8	93.1
STTrack [22]	63.3	63.4	63.2	77.6	82.5	93.7
SUTrack [6]	65.1	65.7	64.5	76.5	82.8	91.8
OneTracker [20]	60.9	60.4	60.7	72.7	81.9	87.2
SDSTrack [21]	61.4	60.9	61.9	72.8	81.2	88.3
UnTrack [56]	61.0	60.8	61.1	72.1	82.0	86.9
ViPT [67]	59.4	59.6	59.2	72.1	81.5	87.1
ProTrack [61]	57.8	57.3	58.3	65.1	80.1	80.2
MCITrack [25]	53.8	54.9	52.7	74.3	82.3	90.6
SPT [69]	53.8	54.9	52.7	65.1	79.8	85.1
SBT-RGBD [58]	-	-	-	70.8	80.9	86.4
OSTrack [62]	52.9	52.2	53.6	67.6	80.3	83.3
DeT [60]	53.2	50.6	56.0	65.7	76.0	84.5
DMTrack [28]	-	-	-	65.8	75.8	85.1
DDiMP [27]	48.5	56.9	50.3	-	-	-
ATCAIS [27]	47.6	45.5	50.0	55.9	76.1	73.9
LTMU-B [10]	46.0	41.7	51.2	-	-	-
GLGS-D [27]	45.3	36.9	58.4	-	-	-
DAL [42]	42.9	36.9	51.2	-	-	-
LTDSEd [26]	40.5	38.2	43.0	-	-	-
Siam-LTD [27]	37.6	34.2	41.8	-	-	-
SiamM-Ds [26]	33.6	26.4	46.3	-	-	-
CA3DMS [35]	22.3	22.8	21.8	-	-	-
DiMP [2]	-	-	-	54.3	70.3	73.1
ATOM [11]	-	-	-	50.5	59.8	68.8

ing modality tracking. Instead, it only incorporates a novel video-level masking strategy. As a result, our tracker outperforms IPT [36] by 2.4% on the LasHeR [31] dataset and by 2.1% on the RGBT234 [29] dataset in terms of P.

RGB-Depth missing modality datasets: We utilize DepthTrack [60] to evaluate the missing modalities in RGB-Depth tracking. The DepthTrack [60] dataset presents various challenges, such as long-term tracking and similar object interference. In long-term videos, video information plays a crucial role in capturing the variations of the object to be tracked. However, once a modality is missing, some trackers fail to track the object due to the lack of long-term contextual information. However, with our masking strategy, our tracker can still benefit from the incomplete modalities in the past frames. Additionally, our tracker employs HMoE-Fuse to handle different cases of missing modalities, enhancing the tracker’s flexibility in dealing with various missing modality scenarios. As shown in Tab. 5, our tracker significantly outperforms other trackers.

4.4. Ablation Study

In subsequent experiments, we conducted evaluations on LasHeR [31] and VisEvent [51] datasets along with their missing-modality variants, using P as the primary metric to assess different methods.

Video-level masking strategy: To enhance the robustness

Table 5. SOTA comparison on DepthTrack_{miss} [60] for RGB-Depth, LasHeR_{miss} [31] and RGBT234_{miss} [29] for RGB-Thermal, and VisEvent_{miss} [51] for RGB-Event.

Method	DepthTrack _{miss} [60]			LasHeR _{miss} [31]		RGBT234 _{miss} [29]		VisEvent _{miss} [51]	
	F-score(↑)	Re(↑)	Pr(↑)	P(↑)	AUC(↑)	MPR(↑)	MSR(↑)	P(↑)	AUC(↑)
FlexTrack	57.8	56.1	59.6	65.1	52.3	84.1	62.6	72.8	55.0
STTrack [22]	49.9	48.8	51.0	54.5	44.9	73.8	54.2	65.5	49.7
SUTrack [6]	49.5	47.3	51.9	58.3	47.6	82.0	60.8	66.6	50.5
SeqTrackv2 [5]	45.0	40.9	50.0	50.0	39.9	70.8	49.9	57.6	43.1
SDSTrack [21]	46.7	42.0	52.7	52.5	43.1	67.0	48.8	62.6	46.9
IPT [36]	-	-	-	61.7	46.4	82.0	59.4	-	-
ViPT [67]	44.4	40.5	46.6	40.1	34.0	52.4	39.4	57.2	43.2
MCITrack [25]	49.7	42.9	59.1	34.2	40.0	53.6	40.9	49.9	36.5

Table 6. Performance comparison of different masking strategies.

	LasHeR	VisEvent	LasHeR _{miss}	VisEvent _{miss}
Ours	77.3	81.4	65.1	72.8
w/o	77.1	81.5	61.3	66.3
Random	77.0	80.5	62.4	70.8
VideoMAE	74.3	78.2	62.8	71.0

of our tracker, we propose a novel video-level masking strategy. To validate its effectiveness, we conducted experiments with and without the masking strategy, as shown in Table 6. For common multimodal datasets (e.g., LasHeR [31] and VisEvent [51]), the performance remains consistent when the masking strategy is removed. However, in scenarios with missing modalities (e.g., LasHeR_{miss} and VisEvent_{miss}), our masking strategy significantly improves performance compared to the baseline without masking. Specifically, we achieve a Precision (P) of 65.1 on LasHeR_{miss} and 72.8 on VisEvent_{miss}, outperforming the no-masking variant by 3.8% and 6.5%, respectively.

Additionally, we compared our strategy with widely used methods, including random masking and VideoMAE [45]. While these methods slightly outperform the no-masking baseline in missing modality scenarios, they lead to performance drops in complete datasets (e.g., LasHeR and VisEvent). For instance, random masking reduces the Precision on LasHeR from 77.3 to 77.0, and VideoMAE further degrades it to 74.3. This suggests that these methods disrupt the temporal continuity required for robust video tracking. In contrast, our video-level masking strategy preserves temporal context, ensuring stable performance across both complete and missing modality scenarios.

HMoE-Fuse The HMoE-Fuse is a crucial component for the efficient fusion of diverse modalities. In Table 7, we present an ablation study examining the impact of various design choices on the performance of HMoE-Fuse. Specifically, "w/o gating" indicates the removal of the gating function, "w/o linear attn." refers to the absence of linear attention, and "w/o HMoE" denotes the removal of the HMoE

Table 7. Ablation studies on HMoE-Fuse.

	LasHeR	VisEvent	LasHeR _{miss}	VisEvent _{miss}
Ours	77.3	81.4	65.1	72.8
w/o gating	76.3	79.9	63.1	71.7
w/o linear attn.	75.4	78.2	61.8	69.9
w/o HMoE	74.2	79.3	60.5	67.5
DeepseekMoE	75.0	81.2	62.1	68.7

Table 8. Ablation study on varying complexity.

	LasHeR	VisEvent	LasHeR _{miss}	VisEvent _{miss}
Ours	77.3	81.4	65.1	72.8
512	76.1	80.9	62.7	71.7
64	76.4	80.1	63.8	70.6
4	77.0	80.8	64.1	71.4

module entirely. The experimental results demonstrate that eliminating these components leads to a significant decline in performance, highlighting their importance. To further validate the effectiveness of our HMoE, we replaced it with DeepSeekMoE [9]. The results show that, for video-level tasks, HMoE-Fuse, which incorporates temporal-spatial information from video clips, outperforms DeepSeekMoE. This is attributed to the heterogeneous experts in HMoE-Fuse, which enables the tracker to have richer representations and better handle the complexities of video data. The superior performance of HMoE-Fuse underscores its advantage in leveraging diverse expert sizes to enhance tracking accuracy and robustness.

In this paper, we employ the HMoE to achieve a more complex and nuanced representation. To validate the effectiveness of our proposed HMoE approach, we conducted experiments using both medium-sized and large-sized experts. As shown in Table 8, our HMoE consistently outperforms other configurations across various metrics. Compared to the homogeneous MoE models, the heterogeneous design of our MoE enables the model to achieve stronger expressive capabilities. This advantage is particularly evident under different missing rates and in challenging scenar-

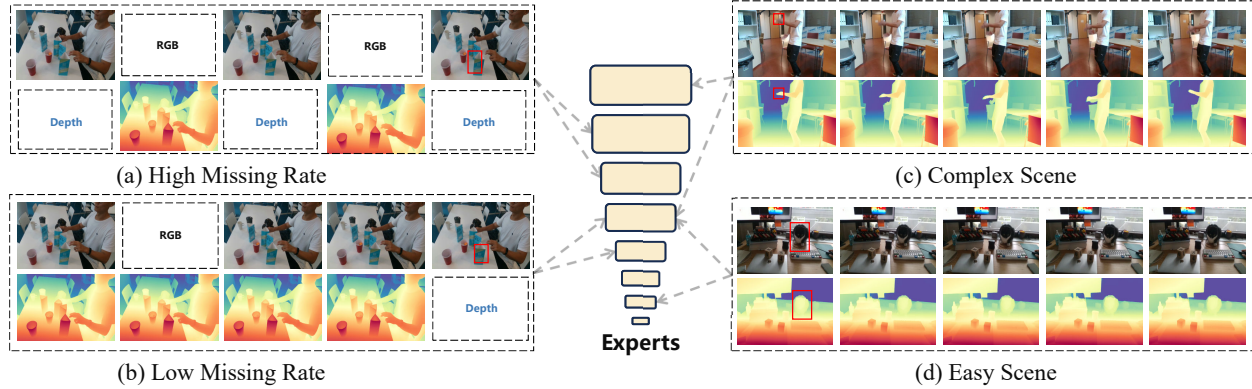


Figure 4. We are interested in the **following question**: *With a higher missing rate, would a model favor a **lighter** expert or a **more complex** expert?* To investigate this, we analyze the impact of the missing modality rate on the same video clip, as shown in (a) and (b). When the missing rate increases (e.g., 50% vs. 20%), we observe that the model adaptively selects more complex experts with larger capacities to compensate for information scarcity through deeper reasoning. This insight leads us to a **second question**: *Does this **behavior** persist in a full-modality setting?* To answer this, we conduct an analysis in (c) and (d), visualizing expert selection in a similar indoor scene under varying levels of difficulty—such as object size (e.g., a moving hand vs. a stationary cat). Our findings reveal that the model prioritizes high-capacity experts to enable fine-grained cross-modal fusion, effectively addressing the challenges posed by dynamic variations in the scene. The block size in the middle represents the hidden dimension of heterogeneous experts in the MoE architecture, where a larger dimension indicates higher complexity.

ios, where our model demonstrates superior performance.

Visualization and Analysis: To gain a deeper understanding of the behavior of our HMoE-Fuse module under varying degrees of missing modalities and full modalities, we conducted a visualization analysis of the selection of experts with different missing rates and different full modalities scenes. The results are presented in Fig. 4.

From the visualizations, we observed that as the missing modality ratio increases, the model tends to select experts with larger hidden dimensions. This is primarily because, when a high proportion of modalities are missing, the model opts for larger experts to ensure precision through more complex reasoning. In contrast, with only a few modalities missing, the model exhibits higher confidence in its predictions and thus favors smaller experts for inference.

A similar pattern is evident in the full-modality scenario under complex conditions. For instance, in Fig. 4 (c), the depth and feature differences between the two hands are minimal, making it challenging for the model to derive accurate tracking results. As a result, the model selects a larger expert to tackle the increased difficulty. Conversely, in Fig. 4 (d), the distinct features of the “cat” compared to its surroundings enable the model to easily pinpoint its location, allowing for the use of smaller experts during tracking.

Takeaways: Our HMoE framework offers a unique approach to interpretable multimodal fusion by explicitly linking the missing modality problem to complexity modeling — where *more modalities reduce complexity*, and vice versa. We observe that, in simpler scenarios when more information is available, the model favors lighter experts, as the increased joint information entropy from the input en-

hances scene understanding. Our framework demonstrates that multimodality not only improves performance but also contributes to a more adaptive system. This insight is introduced and discussed for the first time in the context of multimodal video tracking, aligning with the fundamental principles of multimodal fusion.

5. Conclusion and Future work

In this work, we tackle the critical yet often understudied challenge of robust multimodal tracking under missing modality conditions. To address this, we propose a novel MoE design with heterogeneous experts, paired with a video-level mask training strategy. Additionally, our model can adaptively select experts of different scales during testing, depending on the extent of missing modalities. Experimental results show that our model excels across all leading benchmarks, with both complete and missing modalities.

While FlexTrack enhances multimodal tracking robustness, our future work will focus on integrating generative models for modality reconstruction and exploring lightweight deployment. Nonetheless, this work highlights the potential of heterogeneous expert systems in addressing the inherent unpredictability of real-world multimodal data.

Acknowledgements. The authors sincerely thank the reviewers and all members of the program committee for their tremendous efforts and incisive feedback. This research was supported in part by the Alexander von Humboldt Foundation and in part by the Ministry of Education and Science of Bulgaria (support for INSAIT, under the Bulgarian National Roadmap for Research Infrastructure). C. Ma was supported in part by NSFC (62376156, 62322113).

References

- [1] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *TPAMI*, 41(2):423–443, 2018. 2
- [2] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *ICCV*, 2019. 6
- [3] Yidong Cai, Jie Liu, Jie Tang, and Gangshan Wu. Robust object modeling for visual tracking. In *ICCV*, 2023. 1
- [4] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *CVPR*, 2021. 5
- [5] Xin Chen, Houwen Peng, Dong Wang, Huchuan Lu, and Han Hu. Seqtrack: Sequence to sequence learning for visual object tracking. In *CVPR*, 2023. 1, 2, 7
- [6] Xin Chen, Ben Kang, Wanting Geng, Jiawen Zhu, Yi Liu, Dong Wang, and Huchuan Lu. Sutrack: Towards simple and unified single object tracking. *arXiv preprint arXiv:2412.19138*, 2024. 2, 5, 6, 7
- [7] Zedu Chen, Bineng Zhong, Guorong Li, Shengping Zhang, Ronrong Ji, Zhenjun Tang, and Xianxian Li. Siamban: Target-aware tracking with siamese box adaptive network. *TPAMI*, 45(4):5158–5173, 2022. 5
- [8] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *TPAMI*, 45(9):10850–10869, 2023. 2
- [9] Damai Dai, Chengqi Deng, Chenggang Zhao, RX Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y Wu, et al. Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models. *arXiv preprint arXiv:2401.06066*, 2024. 2, 3, 7
- [10] Kenan Dai, Yunhua Zhang, Dong Wang, Jianhua Li, Huchuan Lu, and Xiaoyun Yang. High-performance long-term tracking with meta-updater. In *CVPR*, 2020. 5, 6
- [11] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Atom: Accurate tracking by overlap maximization. In *CVPR*, 2019. 5, 6
- [12] Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *CVPR*, 2020. 5
- [13] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. Lasot: A high-quality benchmark for large-scale single object tracking. In *CVPR*, 2019. 5
- [14] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *JMLR*, 23(120):1–39, 2022. 5
- [15] Zhihong Fu, Zehua Fu, Qingjie Liu, Wenrui Cai, and Yunhong Wang. Sparset: Visual tracking with sparse transformers. In *IJCAI*, 2022. 1
- [16] Shenyan Gao, Chunlun Zhou, Chao Ma, Xinggang Wang, and Junsong Yuan. Aiatrack: Attention in attention for transformer visual tracking. In *ECCV*. Springer, 2022. 2
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 2
- [18] Dongyan Guo, Jun Wang, Ying Cui, Zhenhua Wang, and Shengyong Chen. Siamcar: Siamese fully convolutional classification and regression for visual tracking. In *CVPR*, 2020. 2
- [19] Xing Han, Huy Nguyen, Carl Harris, Nhat Ho, and Suchi Saria. Fusemoe: Mixture-of-experts transformers for flexible modal fusion. *NeurIPS*, 37:67850–67900, 2025. 2, 3
- [20] Lingyi Hong, Shilin Yan, Renrui Zhang, Wanyun Li, Xinyu Zhou, Pinxue Guo, Kaixun Jiang, Yiting Chen, Jinglun Li, Zhaoyu Chen, et al. Onetracker: Unifying visual object tracking with foundation models and efficient tuning. In *CVPR*, 2024. 2, 5, 6
- [21] Xiaojun Hou, Jiazheng Xing, Yijie Qian, Yaowei Guo, Shuo Xin, Junhao Chen, Kai Tang, Mengmeng Wang, Zhengkai Jiang, Liang Liu, et al. Sdstrack: Self-distillation symmetric adapter learning for multi-modal visual object tracking. In *CVPR*, 2024. 2, 5, 6, 7
- [22] Xiantao Hu, Ying Tai, Xu Zhao, Chen Zhao, Zhenyu Zhang, Jun Li, Bineng Zhong, and Jian Yang. Exploiting multimodal spatial-temporal patterns for video object tracking. *arXiv preprint arXiv:2412.15691*, 2024. 2, 5, 6, 7
- [23] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *TPAMI*, 43(5):1562–1577, 2019. 5
- [24] Lan Jiang, Ye Mao, Xiangfeng Wang, Xi Chen, and Chao Li. Cola-diff: Conditional latent diffusion model for multimodal mri synthesis. In *MICCAI*, pages 398–408. Springer, 2023. 2
- [25] Ben Kang, Xin Chen, Simiao Lai, Yang Liu, Yi Liu, and Dong Wang. Exploring enhanced contextual information for video-level object tracking. *arXiv preprint arXiv:2412.11023*, 2024. 5, 6, 7
- [26] Matej Kristan, Jiri Matas, Ales Leonardis, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kamarainen, Luka Čehovin Zajc, Ondrej Drbohlav, Alan Lukežič, Amanda Berg, et al. The seventh visual object tracking VOT2019 challenge results. In *ICCVW*, 2019. 6
- [27] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, Ondrej Drbohlav, et al. The eighth visual object tracking VOT2020 challenge results. In *ECCVW*, 2020. 6
- [28] Matej Kristan, Aleš Leonardis, Jiří Matas, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Hyung Jin Chang, Martin Danelljan, Luka Čehovin Zajc, Alan Lukežič, et al. The tenth visual object tracking vot2022 challenge results. In *ECCVW*, 2023. 6
- [29] Chenglong Li, Xinyan Liang, Yijuan Lu, Nan Zhao, and Jin Tang. Rgb-t object tracking: Benchmark and baseline. *PR*, 96:106977, 2019. 1, 2, 5, 6, 7
- [30] Chenglong Li, Lei Liu, Andong Lu, Qing Ji, and Jin Tang. Challenge-aware RGBT tracking. In *ECCV*, 2020. 6
- [31] Chenglong Li, Wanlin Xue, Yaqing Jia, Zhichen Qu, Bin Luo, Jin Tang, and Dengdi Sun. Lasher: A large-scale high-diversity benchmark for rgbt tracking. *TIP*, 31:392–404, 2021. 1, 5, 6, 7
- [32] Wei Li, Zengfu Hou, Jun Zhou, and Ran Tao. Siambag: Band attention grouping-based siamese object tracking network for hyperspectral videos. *TGRS*, 61:1–12, 2023. 2

- [33] Liting Lin, Heng Fan, Zhipeng Zhang, Yong Xu, and Haibin Ling. Swintrack: A simple and strong baseline for transformer tracking. *NeurIPS*, 35:16743–16754, 2022. 1, 2
- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5
- [35] Ye Liu, Xiao-Yuan Jing, Jianhui Nie, Hao Gao, Jun Liu, and Guo-Ping Jiang. Context-aware three-dimensional mean-shift with occlusion handling for robust object tracking in RGB-D videos. *TMM*, pages 664–677, 2018. 6
- [36] Andong Lu, Chenglong Li, Jiacong Zhao, Jin Tang, and Bin Luo. Modality-missing rgbt tracking: Invertible prompt learning and high-quality benchmarks. *IJCV*, pages 1–21, 2024. 1, 2, 6, 7
- [37] Mengmeng Ma, Jian Ren, Long Zhao, Sergey Tulyakov, Cathy Wu, and Xi Peng. Smil: Multimodal learning with severely missing modality. In *AAAI*, pages 2302–2310, 2021. 1, 2
- [38] Mengmeng Ma, Jian Ren, Long Zhao, Davide Testuggine, and Xi Peng. Are multimodal transformers robust to missing modality? In *CVPR*, pages 18177–18186, 2022. 1, 2
- [39] Christoph Mayer, Martin Danelljan, Danda Pani Paudel, and Luc Van Gool. Learning target candidate association to keep track of what not to track. In *ICCV*, 2021. 1
- [40] Matthias Muller, Adel Bibi, Silvio Giancola, Salman Alsubaihi, and Bernard Ghanem. Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In *ECCV*, 2018. 5
- [41] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016. 5
- [42] Yanlin Qian, Song Yan, Alan Lukežič, Matej Kristan, Joni-Kristian Kämäräinen, and Jiri Matas. Dal: A deep depth-aware long-term tracker. In *ICPR*, 2021. 6
- [43] Yibing Song, Chao Ma, Xiaohu Wu, Lijun Gong, Linchao Bao, Wangmeng Zuo, Chunhua Shen, Rynson WH Lau, and Ming-Hsuan Yang. Vital: Visual tracking via adversarial learning. In *CVPR*, 2018. 5
- [44] Yunjie Tian, Lingxi Xie, Jihao Qiu, Jianbin Jiao, Yaowei Wang, Qi Tian, and Qixiang Ye. Fast-itpn: Integrally pre-trained transformer pyramid network with token migration. *TPAMI*, 2024. 5
- [45] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *NeurIPS*, 35:10078–10093, 2022. 4, 7
- [46] Paul Voigtlaender, Jonathon Luiten, Philip H. S. Torr, and Bastian Leibe. Siam R-CNN: Visual tracking by re-detection. In *CVPR*, 2020. 5
- [47] Chaoqun Wang, Chunyan Xu, Zhen Cui, Ling Zhou, Tong Zhang, Xiaoya Zhang, and Jian Yang. Cross-modal pattern-propagation for rgb-t tracking. In *CVPR*, 2020. 6
- [48] Hu Wang, Congbo Ma, Jianpeng Zhang, Yuan Zhang, Jodie Avery, Louise Hull, and Gustavo Carneiro. Learnable cross-modal knowledge distillation for multi-modal learning with missing modality. In *MICCAI*, pages 216–226. Springer, 2023. 2
- [49] Qiang Wang, Li Zhang, Luca Bertinetto, Weiming Hu, and Philip H. S. Torr. Fast online object tracking and segmentation: A unifying approach. In *CVPR*, 2019. 5
- [50] Xiao Wang, Jianing Li, Lin Zhu, Zhipeng Zhang, Zhe Chen, Xin Li, Yaowei Wang, Yonghong Tian, and Feng Wu. Visevent: Reliable object tracking via collaboration of frame and event flows. *IEEE Transactions on Cybernetics*, 2023. 1
- [51] Xiao Wang, Jianing Li, Lin Zhu, Zhipeng Zhang, Zhe Chen, Xin Li, Yaowei Wang, Yonghong Tian, and Feng Wu. Visevent: Reliable object tracking via collaboration of frame and event flows. *TCYB*, pages 1–14, 2023. 2, 5, 6, 7
- [52] Xiao Wang, Shiao Wang, Chuanming Tang, Lin Zhu, Bo Jiang, Yonghong Tian, and Jin Tang. Event stream-based visual object tracking: A high-resolution benchmark dataset and a novel baseline. In *CVPR*, 2024. 1
- [53] Xing Wei, Yifan Bai, Yongchao Zheng, Dahu Shi, and Yihong Gong. Autoregressive visual tracking. In *CVPR*, 2023. 1, 2
- [54] Sangmin Woo, Sumin Lee, Yeonju Park, Muhammad Adi Nugroho, and Changick Kim. Towards good practices for missing modality robust action recognition. In *AAAI*, pages 2776–2784, 2023. 2
- [55] Qiangqiang Wu, Tianyu Yang, Ziquan Liu, Baoyuan Wu, Ying Shan, and Antoni B Chan. Dropmae: Masked autoencoders with spatial-attention dropout for tracking tasks. In *CVPR*, 2023. 4
- [56] Zongwei Wu, Jilai Zheng, Xiangxuan Ren, Florin-Alexandru Vasluianu, Chao Ma, Danda Pani Paudel, Luc Van Gool, and Radu Timofte. Single-model and any-modality for video object tracking. In *CVPR*, 2024. 2, 5, 6
- [57] Yun Xiao, Mengmeng Yang, Chenglong Li, Lei Liu, and Jin Tang. Attribute-based progressive fusion network for rgbt tracking. In *AAAI*, 2022. 6
- [58] Fei Xie, Chunyu Wang, Guangting Wang, Yue Cao, Wankou Yang, and Wenjun Zeng. Correlation-aware deep tracking. In *CVPR*, 2022. 6
- [59] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *ICCV*, 2021. 5
- [60] Song Yan, Jinyu Yang, Jani Käpylä, Feng Zheng, Aleš Leonardis, and Joni-Kristian Kämäräinen. Depthtrack: Unveiling the power of rgb-d tracking. In *ICCV*, 2021. 1, 5, 6, 7
- [61] Jinyu Yang, Zhe Li, Feng Zheng, Ales Leonardis, and Jingkuan Song. Prompting for multi-modal tracking. In *ACMMM*, 2022. 5, 6
- [62] Botao Ye, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Joint feature learning and relation modeling for tracking: A one-stream framework. In *ECCV*, 2022. 2, 5, 6
- [63] Sukwon Yun, Inyoung Choi, Jie Peng, Yangfan Wu, Jingxuan Bao, Qiyiwen Zhang, Jiayi Xin, Qi Long, and Tianlong Chen. Flex-moe: Modeling arbitrary modality combination via the flexible mixture-of-experts. In *NeurIPS*. 2
- [64] Hui Zhang, Lei Zhang, Li Zhuo, and Jing Zhang. Object tracking in RGB-T videos using modal-aware attention network and competitive learning. *Sensors*, page 393, 2020. 6

- [65] Pengyu Zhang, Jie Zhao, Chunjuan Bo, Dong Wang, Huchuan Lu, and Xiaoyun Yang. Jointly modeling motion and appearance cues for robust RGB-T tracking. *TIP*, pages 3335–3347, 2021. [6](#)
- [66] Jinming Zhao, Ruichen Li, and Qin Jin. Missing modality imagination network for emotion recognition with uncertain missing modalities. In *ACL*, pages 2608–2618, 2021. [1](#), [2](#)
- [67] Jiawen Zhu, Simiao Lai, Xin Chen, Dong Wang, and Huchuan Lu. Visual prompt multi-modal tracking. In *CVPR*, 2023. [2](#), [5](#), [6](#), [7](#)
- [68] Tong Zhu, Xiaoye Qu, Daize Dong, Jiacheng Ruan, Jingqi Tong, Conghui He, and Yu Cheng. Llama-moe: Building mixture-of-experts from llama with continual pre-training. *arXiv preprint arXiv:2406.16554*, 2024. [2](#)
- [69] Xue-Feng Zhu, Tianyang Xu, Zhangyong Tang, Zucheng Wu, Haodong Liu, Xiao Yang, Xiao-Jun Wu, and Josef Kittler. Rgbd1k: A large-scale dataset and benchmark for rgb-d object tracking. In *AAAI*, 2023. [2](#), [6](#)
- [70] Yabin Zhu, Chenglong Li, Jin Tang, and Bin Luo. Quality-aware feature aggregation network for robust RGBT tracking. *TIV*, pages 121–130, 2020. [6](#)