

RAGTrack: Language-aware RGBT Tracking with Retrieval-Augmented Generation

Hao Li^{1,2}, Yuhao Wang², Wenning Hao^{1*}, Pingping Zhang^{2*}, Dong Wang³, Huchuan Lu^{2,3}

¹College of Command and Control Engineering, Army Engineering University of PLA

²School of Future Technology, Dalian University of Technology

³School of Information and Communication Engineering, Dalian University of Technology

lihao@aeu.edu.cn, 924973292@mail.dlut.edu.cn, hwnbox@aeu.edu.cn,

{zhpp, wdice, lhchuan}@dlut.edu.cn

Abstract

*RGB-Thermal (RGBT) tracking aims to achieve robust object localization across diverse environmental conditions by fusing visible and thermal infrared modalities. However, existing RGBT trackers rely solely on initial-frame visual information for target modeling, failing to adapt to appearance variations due to the absence of language guidance. Furthermore, current methods suffer from redundant search regions and heterogeneous modality gaps, causing background distraction. To address these issues, we first introduce textual descriptions into RGBT tracking benchmarks. This is accomplished through a pipeline that leverages Multi-modal Large Language Models (MLLMs) to automatically produce textual annotations. Afterwards, we propose **RAGTrack**, a novel **R**etrieval-**A**ugmented **G**eneration framework for robust RGBT tracking. To this end, we introduce a Multi-modal Transformer Encoder (MTE) for unified visual-language modeling. Then, we design an Adaptive Token Fusion (ATF) to select target-relevant tokens and perform channel exchanges based on cross-modal correlations, mitigating search redundancies and modality gaps. Finally, we propose a Context-aware Reasoning Module (CRM) to maintain a dynamic knowledge base and employ a Retrieval-Augmented Generation (RAG) to enable temporal linguistic reasoning for robust target modeling. Extensive experiments on four RGBT benchmarks demonstrate that our framework achieves state-of-the-art performance across various challenging scenarios. The source code is available at <https://github.com/IdolLab/RAGTrack>.*

1. Introduction

Visual object tracking [6–10, 29–35, 47, 50–57, 76–81] aims to localize objects of interest across video sequences.

*Corresponding author.

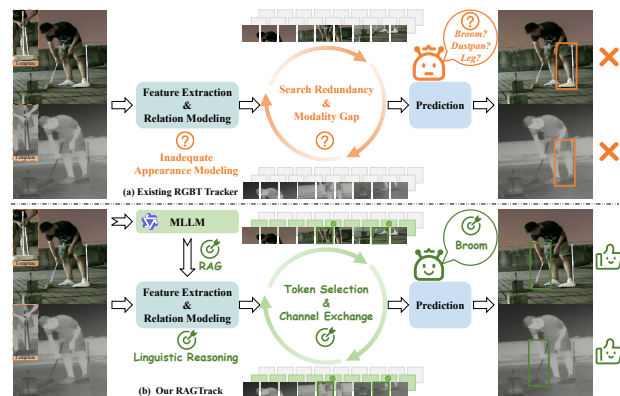


Figure 1. Comparison with different RGBT tracking paradigms. (a) Existing RGBT trackers suffer from inadequate appearance modeling, search redundancy, and modality gap. (b) Our RAGTrack introduces linguistic reasoning, dynamic token selection, and adaptive channel exchange.

Despite its wide applications in autonomous driving, intelligent surveillance, and human-computer interaction, tracking relying solely on the visible (RGB) modality often fails under low illumination. In contrast, the thermal infrared (TIR) imaging has the advantage of all-weather capability, providing complementary information of the RGB modality. To integrate these advantages, RGB-Thermal (RGBT) tracking combines both modalities for robust performance.

However, current RGBT trackers [28, 38–43] model the target using visual information from the initial frame, which often leads to drift under severe appearance variations. This limitation stems from two main factors. First, a single image template offers limited information of the target, failing to capture full appearance changes across different view-points. Second, the inherent ambiguity of targets often causes trackers to overemphasize unstable features, resulting in background distraction. As shown in Fig. 1, the tracker may become confused among the broom, dustpan, or the lower body of the pedestrian. Moreover, bounding

box-based initialization is inconvenient in practical scenarios, consequently hindering its widespread deployment.

On the other hand, recent advances in multi-modal object tracking have focused on integrating features from diverse modalities. In previous works, multi-modal features are first extracted independently, then fused through dedicated fusion modules [63]. They establish connections of multiple feature sources, including RGB and TIR templates and search regions, operating at token-level. However, a significant portion of search regions is redundant, containing substantial background regions and distractors that negatively impact the tracking accuracy. Furthermore, heterogeneous modality gaps hinder the establishment of effective cross-modality correspondences. This raises a basic question: *How can we construct a target model with rich semantic representations while mitigating the impact of search redundancies and modality gaps?*

To address these challenges, we introduce language as a high-level representation to enhance RGBT tracking. Language provides a more abstract understanding of objects than images, addressing semantic limitations in visual representation. Language-based target representation can offer clear semantic information, including object categories, appearance attributes and motion states, enabling effective target-background separation. Recent advances in RGB-Language (RGL) tracking have demonstrated the potential of language descriptions for robust tracking [9]. However, these methods face the challenge of visual-language misalignment as targets evolve across video frames. Recently, Multi-modal Large Language Models (MLLMs) have exhibited strong reasoning capabilities, delivering opportunities for incorporating textual descriptions into RGBT tracking. Despite these advancements, there is no RGBT tracking benchmarks with text annotations. To bridge this gap, we extend current RGBT tracking benchmarks by introducing textual descriptions through MLLMs.

Building upon the textual descriptions, we propose a novel RGBT tracking framework, named RAGTrack. As depicted in Fig. 2, our framework consists of three key components: Multi-modal Transformer Encoder (MTE), Adaptive Token Fusion (ATF) and Context-aware Reasoning Module (CRM). The MTE utilizes text features for unified modeling with visual patch tokens. Meanwhile, ATF utilizes linguistic cues to obtain text-guided attention scores for identifying target-relevant search tokens, reducing noisy interference. This module further employs adaptive channel exchange to dynamically bridge heterogeneous features across modalities. As a result, it maintains semantic consistency and mitigates modality gaps. The CRM implements a Retrieval-Augmented Generation (RAG) paradigm that establishes a continuous reasoning cycle across frames. This module maintains a dynamic knowledge base, retrieves historically relevant features, and generates adaptive tar-

get descriptions via MLLMs. With the above components, our RAGTrack enables comprehensive temporal reasoning, leveraging both visual evidence and linguistic context to maintain target identities under complex scenarios. Extensive experiments on four benchmarks demonstrate the effectiveness of our method, achieving state-of-the-art performance. In summary, our main contributions are as follows:

- To the best of our knowledge, we are the first to introduce textual descriptions into RGBT tracking, extending existing benchmarks with semantic annotations.
- We propose RAGTrack, a novel framework that leverages Retrieval-Augmented Generation (RAG) to enhance linguistic reasoning for robust target modeling.
- We design the Adaptive Token Fusion (ATF) to address search redundancies and modality gaps through dynamic token selection and adaptive channel exchange.
- Extensive experiments on four RGBT tracking benchmarks demonstrate that RAGTrack outperforms existing methods in both accuracy and robustness.

2. Related Work

2.1. RGB-Thermal Tracking

RGBT tracking [37] combines RGB and TIR modalities to overcome single-modality limitations in complex scenarios. Early works [85] use MDNet [46] for feature extraction, while CMPP [61] employs affinity-based pattern propagation. Subsequent methods advance fusion strategies, such as template-bridged interactions [17] and Mamba-based fusion [42]. The Mixture of Experts (MoE) paradigms, exemplified by XTrack [57] and FlexTrack [56], further enhance robustness through expert routing and heterogeneous fusion. Meanwhile, temporal modeling is refined by TATrack [62], CStrack [8] and STTrack [15]. Recently, ViPT [83], BAT [2], SDSTrack [14], OneTracker [13] and RDTTrack [84] leverage visual prompt learning to boost tracking performance. However, current methods solely depend on visual information, which becomes unstable during tracking. Different from existing methods, our method leverages textual descriptions to complement visual features and distinguish similar-appearance targets.

2.2. RGB-Language Tracking

RGL tracking [7] utilizes both natural language descriptions and visual references to localize the target across video sequences. Li et al. [32] pioneer this field by extending tracking datasets with linguistic annotations. Wang et al. [65] and Zhou et al. [82] advance visual grounding techniques to integrate language semantics with tracking. Subsequently, SNLT [6], All-in-One [71] and OVLM [72] develop multi-modal fusion approaches through dynamic aggregation, unified Transformers and memory mechanism, respectively. CiteTracker [29] and MMTrack [78]

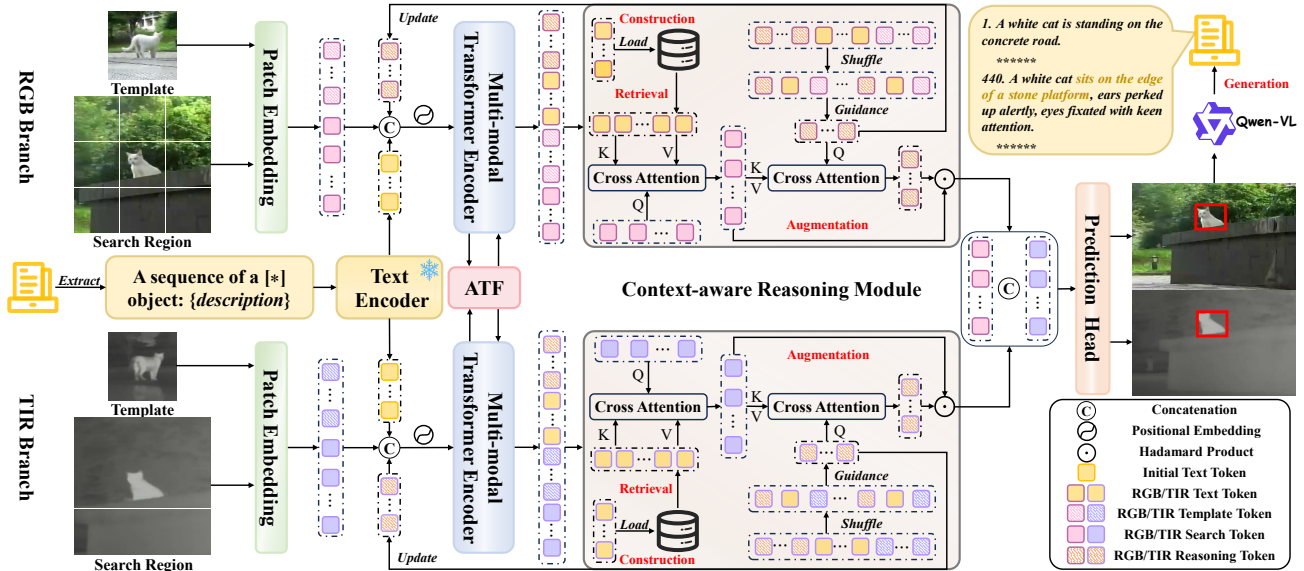


Figure 2. Overall framework. Our method begins by tokenizing input texts and images with reasoning tokens. MTE then performs unified visual-language modeling, while ATF utilizes text-guided attention to dynamically select target-relevant tokens and enables adaptive channel exchange. Subsequently, CRM retrieves relevant contexts from a dynamic knowledge base for context-aware reasoning. Finally, the prediction head outputs tracking results, which are used by MLLMs to generate updated textual descriptions for following frames.

adopt generative frameworks for text-based target estimation, while VLTTT [12] and UVLTrack [44] leverage contrastive learning for vision-language alignment. Recently, LLM-driven methods, such as DTLLM-VLT [30], ChatTracker [55], DUTrack [31] and ReasoningTrack [66], enhance robustness through multimodal reasoning and dynamic language generation. However, current methods suffer from insufficient visual-language feature association and ineffective background interference suppression. Different from existing methods, our method utilizes text features to address search redundancies and modality gaps through dynamic token selection and adaptive channel exchange.

2.3. Retrieval-Augmented Generation

RAG [21] can enhance LLMs by retrieving external knowledge during inference to mitigate hallucinations and improve the factual accuracy in domain-specific tasks. While widely adopted in applications [11, 58], its integration into visual tracking remains limited. As an outstanding work, TrackingMiM [33] proposes a query-based retrieval module for UAV tracking. However, current methods merely reuse pre-stored features. Different from existing methods, our method introduces RAG into RGBT tracking for the first time, enabling context-aware linguistic reasoning.

3. Methodology

In this paper, we propose RAGTrack, which includes three key components: Multi-modal Transformer Encoder (MTE), Adaptive Token Fusion (ATF), and Context-aware Reasoning Module (CRM). As shown in Fig. 2, MTE serves

as the foundation for processing heterogeneous inputs. ATF dynamically captures target-relevant tokens and fuses cross-modal features. CRM leverages a RAG mechanism to maintain target identities through temporal semantic reasoning. Details are described as follows.

3.1. Overall Framework

Our framework localizes targets across video sequences by fusing complementary information from RGB, TIR, and language modalities. At time step t , the input consists of search images $\mathbf{X}_m^t \in \mathbb{R}^{3 \times H_x \times W_x}$ and multi-modal references comprising template images $\mathbf{Z}_m^t \in \mathbb{R}^{3 \times H_z \times W_z}$ and language descriptions \mathbf{L}^t . Here, H and W denote the height and width of the images, and $m \in \{B, R\}$ represents the RGB or TIR modality. The MTE employs parameter-shared branches for processing both RGB and TIR modalities. The encoder processes these inputs via patch embedding and tokenization to generate tokens, which are concatenated and fed into the visual-language unified modeling. These tokens then undergo ATF to mitigate search redundancies and modality gaps, followed by CRM for linguistic refinement. Finally, the refined features are fed into the prediction head to produce tracking results.

3.2. Multi-modal Transformer Encoder

To extract multi-modal features, we introduce MTE for unified visual-language modeling. We process template and search images through a three-stage downsampling [75]. It transforms the template and search images into patch tokens $\hat{\mathbf{Z}}_m^t \in \mathbb{R}^{N_z \times C}$ and $\hat{\mathbf{X}}_m^t \in \mathbb{R}^{N_x \times C}$, where N represents the




Figure 3. Details of our proposed ATF.

number of tokens, and C represents the feature dimension.

Afterwards, we introduce a sequence prefix \mathbf{E}^t to enhance the temporal awareness and address the potential misalignment between visual contents and language descriptions across frames. This prefix combines a fixed textual prompt with learnable tokens. Specifically, \mathbf{E}^t is formulated as “A sequence of a [*] object:”, where [*] is replaced by learnable tokens. We concatenate \mathbf{E}^t and \mathbf{L}^t to form the text input $\mathbf{H}^t = [\mathbf{E}^t, \mathbf{L}^t]$. Then, \mathbf{H}^t is encoded by the text encoder \mathcal{T} to obtain semantic features $\hat{\mathbf{H}}^t \in \mathbb{R}^{N_h \times C}$:

$$\hat{\mathbf{H}}^t = \mathcal{T}(\mathbf{H}^t). \quad (1)$$

Subsequently, we concatenate the reasoning token $\mathbf{R}_m^t \in \mathbb{R}^{N_r \times C}$, the text token $\hat{\mathbf{H}}^t$, the template token $\hat{\mathbf{Z}}_m^t$ and the search token $\hat{\mathbf{X}}_m^t$ into a unified token sequence:

$$\mathbf{F}_m^0 = \left[\mathbf{R}_m^t; \hat{\mathbf{H}}^t; \hat{\mathbf{Z}}_m^t; \hat{\mathbf{X}}_m^t \right], \quad (2)$$

where $[\cdot]$ denotes the concatenation operation. The unified visual-language modeling is formulated as:

$$\begin{aligned} \hat{\mathbf{F}}_m^{l-1} &= \text{MHSA}(\mathbf{F}_m^{l-1}, \mathbf{F}_m^{l-1}, \mathbf{F}_m^{l-1}), \\ \tilde{\mathbf{F}}_m^{l-1} &= \mathbf{F}_m^{l-1} + \text{LN}(\delta_1 \cdot \hat{\mathbf{F}}_m^{l-1}), \\ \mathbf{F}_m^l &= \tilde{\mathbf{F}}_m^{l-1} + \text{LN}(\delta_2 \cdot \text{MLP}(\tilde{\mathbf{F}}_m^{l-1})), \end{aligned} \quad (3)$$

where MHSA stands for multi-head self-attention [5], LN represents the layer normalization [1] and MLP denotes the multilayer perceptron. $l \in \{1, 2, \dots, L\}$ indicates the layer index. δ_1 and δ_2 are two learnable parameters. By unified visual-language modeling, semantic information from text is effectively leveraged to enhance feature discrimination.

3.3. Adaptive Token Fusion

To address search redundancies and modality gaps, we propose ATF for dynamic token selection and adaptive channel exchange. Traditional fusion methods suffer from inefficient token utilization and significant feature discrepancies, leading to degraded tracking performance. In contrast, ATF identifies target-relevant tokens and adapts channel exchanges based on cross-modality correlations.

In fact, attention scores \mathbf{A}_m from the self-attention module serve as an effective indicator for token significance, while simultaneously guiding the model to excavate contextual information. This fact motivates our use of attention

scores as metrics for token selection. As shown in Fig. 3, the calculation for each term is as follows:

$$\mathbf{A}_m^{x2o} = \text{Softmax} \left(\frac{\mathbf{Q}_m^x (\mathbf{K}_m^o)^T}{\sqrt{d}} \right), \quad (4)$$

where \mathbf{A}_m^{x2r} , \mathbf{A}_m^{x2h} , \mathbf{A}_m^{x2z} , and \mathbf{A}_m^{x2x} represent correlations between search and reasoning, text, template, and search tokens, respectively. Specially, to mitigate template noise, we extract the center region of the template image, as it contains sufficient target information [70]. We aggregate these scores to identify key tokens as follows:

$$\mathbf{A}_m^{\text{total}} = \mathbf{A}_m^{x2r} + \mathbf{A}_m^{x2h} + \mathbf{A}_m^{x2z} + \mathbf{A}_m^{x2x}. \quad (5)$$

Tokens with higher $\mathbf{A}_m^{\text{total}}$ are retained using a retention ratio γ . This parameter-free token selection avoids redundant computations by directly reusing the attention scores.

For channel exchange, we address the challenge of modality gaps by establishing channel-level associations between modalities. More specifically, we first compute cross-modal relevance \mathbf{S} between features along the channel dimension:

$$\mathbf{S} = \left((\mathbf{F}_B^l)^T \mathbf{W}_B^l \right) \left((\mathbf{F}_R^l)^T \mathbf{W}_R^l \right)^T, \quad (6)$$

where \mathbf{W}_B^l and \mathbf{W}_R^l are learnable parameters. Then, we average \mathbf{S} along the channel dimension to determine channel importance, then select corresponding channels for adaptive exchange based on the exchange ratio σ . Finally, we concatenate the features of RGB and TIR modalities along the token dimension and apply the cross-modal fusion \mathcal{M} :

$$\mathbf{F}_B^l, \mathbf{F}_R^l = \mathcal{M} \left([\mathbf{F}_B^l; \mathbf{F}_R^l] \right), \quad (7)$$

where \mathcal{M} is a MLP layer. This channel exchange bridges the independently processed RGB and TIR modalities during feature extraction, enabling the model to learn cross-modal correlations and enhance feature discrimination.

Our ATF mitigates search redundancies and modality gaps through dynamic token selection and adaptive channel exchange. It facilitates information transmission between RGB and TIR modalities, enabling the model to focus on the most relevant features for enhanced robustness.

3.4. Context-aware Reasoning Module

To enhance temporal modeling, we propose the CRM for linguistic reasoning. As illustrated in Fig. 2, it includes a RAG mechanism that operates through four stages: construction, retrieval, augmentation, and generation.

Construction. We construct a local knowledge base from historical text features. This knowledge base is maintained as a set of n feature embeddings $\mathbf{D}_m = \{\hat{\mathbf{H}}_m^1, \dots, \hat{\mathbf{H}}_m^n\}$, which is dynamically updated during inference to suppress redundancy. Specifically, a new text

feature $\hat{\mathbf{H}}_m^t$ is added only if its maximum cosine similarity with existing entries falls below a threshold λ :

$$\max_{\hat{\mathbf{H}}_m^i \in \mathbf{D}_m} \frac{\hat{\mathbf{H}}_m^t \cdot \hat{\mathbf{H}}_m^i}{\|\hat{\mathbf{H}}_m^t\|_2 \|\hat{\mathbf{H}}_m^i\|_2} < \lambda. \quad (8)$$

Retrieval. Given a query feature $\hat{\mathbf{H}}_m^t$, the retriever \mathcal{O} selects the top- k most relevant features from \mathbf{D}_m :

$$\mathbf{V}_m = \mathcal{O}(\hat{\mathbf{H}}_m^t, \mathbf{D}_m), \quad (9)$$

where $\mathbf{V}_m \subseteq \mathbf{D}_m$ and $|\mathbf{V}_m| = k$. To refine the search features, we then apply intra-modal cross-attention Φ :

$$\tilde{\mathbf{X}}_m^t = \hat{\mathbf{X}}_m^t + \Phi(\hat{\mathbf{X}}_m^t, \mathbf{V}_m). \quad (10)$$

Augmentation. For temporal propagation across video frames, we enhance the reasoning tokens to capture both local and global context. First, we perform an average pooling \mathcal{P} on the reasoning features, text features, and template features along the token dimension:

$$\bar{\mathbf{R}}_m^t, \bar{\mathbf{H}}_m^t, \bar{\mathbf{Z}}_m^t = \mathcal{P}(\mathbf{R}_m^t, \hat{\mathbf{H}}_m^t, \hat{\mathbf{Z}}_m^t). \quad (11)$$

Subsequently, we concatenate these pooled features along the channel dimension and apply the guidance \mathcal{G} to inject current frame cues:

$$\mathbf{R}_m^{t+1} = \mathcal{G}(\bar{\mathbf{R}}_m^t, \bar{\mathbf{H}}_m^t, \bar{\mathbf{Z}}_m^t), \quad (12)$$

where \mathcal{G} is a MLP layer. Then, the updated reasoning token \mathbf{R}_m^{t+1} is propagated to the next frame.

Finally, we enhance the feature representation through a three-step temporal augmentation:

$$\begin{aligned} \hat{\mathbf{R}}_m^{t+1} &= \mathbf{R}_m^{t+1} + \Phi(\mathbf{R}_m^{t+1}, \bar{\mathbf{X}}_m^t), \\ \tilde{\mathbf{R}}_m^{t+1} &= \hat{\mathbf{R}}_m^{t+1} + \text{MLP}(\hat{\mathbf{R}}_m^{t+1}), \\ \tilde{\mathbf{X}}_m^t &= \bar{\mathbf{X}}_m^t \otimes (\tilde{\mathbf{R}}_m^{t+1})^T \odot \bar{\mathbf{X}}_m^t, \end{aligned} \quad (13)$$

where \otimes and \odot denotes the matrix multiplication and hadamard product.

Generation. To overcome the limitations of static language annotations, we employ MLLMs to dynamically generate context-aware target descriptions during inference. The input consists of the search image and a structured prompt: “Describe the object located in the image at $\langle \text{box} \rangle(x,y,x+w,y+h)\langle / \text{box} \rangle$. Focus on distinctive visual features, motion patterns, and key identifiers to distinguish it from background elements and distractors.” The resulting descriptions continuously refresh the multi-modal reference, improving cross-frame appearance reasoning.

Our CRM maintains a dynamic knowledge base through iterative retrieval and update mechanisms, and enhances temporal coherence via reasoning token propagation. It further generates adaptive target descriptions to overcome static annotation limitations and adapt to scene variations.

3.5. Prediction Head and Loss Function

The refined multi-modal features $\tilde{\mathbf{X}}_m^t$ are first reshaped into a 2D spatial representation. These features are then concatenated, followed by channel compression with convolutional layers. Finally, they are transformed through a Fully Convolutional Network (FCN) consisting of stacked Conv-BN-ReLU layers [18, 45]. The prediction head generates three outputs: classification scores, spatial offsets, and normalized sizes. The final bounding box is constructed at the position with maximum classification score by combining the corresponding predictions.

For model optimization, we adopt a multi-task loss function combining the focal loss L_{cls} [20] for classification, along with L_1 loss and generalized IoU loss L_{iou} [49] for regression. The overall loss function is formulated as:

$$\mathcal{L} = L_{\text{cls}} + \lambda_{\text{iou}} L_{\text{iou}} + \lambda_{L_1} L_1, \quad (14)$$

where λ_{iou} and λ_{L_1} are the hyper-parameters.

4. Experiment

4.1. Datasets and Evaluation Metrics

Our experiments are conducted on four RGBT tracking benchmarks. GTOT [23] contains 50 RGB-T sequences covering seven primary challenges. RGBT210 [24] contains 210 sequences and 12 challenging attributes. RGBT234 [25] includes 234 sequences with increased diversities. LasHeR [27] comprises 1,224 sequences annotated with 19 fine-grained attributes. To enable language-aware RGBT tracking, we extend these benchmarks with textual descriptions through a two-step generation pipeline. In the first step, we generate informative text descriptions by inputting the image and bounding box into MLLMs, overcoming the time-consuming and labor-intensive nature of manual annotations. Then, we refine these generated descriptions through MLLMs and human experts to mitigate hallucinations, ensuring high-quality semantic annotations. For the training set of LasHeR, we annotate all frames across 979 sequences with 514,081 textual descriptions to support model training. The test sets of LasHeR and the other three benchmarks are annotated with first-frame descriptions to evaluate the textual reasoning capability during inference. Following standard protocols, we adopt the Precision Rate (PR) and Success Rate (SR) as primary evaluation metrics. For LasHeR, we additionally use the Normalized Precision Rate (NPR) to account for target scale variations. On GTOT and RGBT234 with modality misalignment, we report the Maximum Precision Rate (MPR) and Maximum Success Rate (MSR) for fair comparison.

4.2. Implementation Details

We implement RAGTrack using PyTorch on 4 NVIDIA V100 GPUs with a batch size of 16. The text encoder em-

Table 1. Performance comparison on four RGBT tracking benchmarks. The best results are in bold.

Method	Publication	Resolution	GTOT		RGBT210		RGBT234		LasHeR		
			MPR \uparrow	MSR \uparrow	PR \uparrow	SR \uparrow	MPR \uparrow	MSR \uparrow	PR \uparrow	NPR \uparrow	SR \uparrow
CMPP [61]	CVPR 2020	256 \times 256	92.6	73.8	—	—	82.3	57.5	—	—	—
CAT [26]	ECCV 2020	256 \times 256	88.9	71.7	79.2	53.3	80.4	56.1	45.0	39.5	31.4
APFNet [68]	AAAI 2022	256 \times 256	90.5	73.7	—	—	82.7	57.9	50.0	43.9	36.2
CMD [73]	CVPR 2023	256 \times 256	89.2	73.4	—	—	82.4	58.4	59.0	54.6	46.4
ViPT [83]	CVPR 2023	256 \times 256	—	—	—	—	83.5	61.7	65.1	—	52.5
TBSI [17]	CVPR 2023	256 \times 256	—	—	85.3	62.5	87.1	63.7	69.2	65.7	55.6
TATrack [62]	AAAI 2024	256 \times 256	—	—	85.3	61.8	87.2	64.4	70.2	66.7	56.1
BAT [2]	AAAI 2024	256 \times 256	—	—	—	—	86.8	64.1	70.2	—	56.3
GMMT [59]	AAAI 2024	256 \times 256	—	—	—	—	87.9	64.7	70.7	67.0	56.6
Un-Track [67]	CVPR 2024	256 \times 256	—	—	—	—	83.7	61.8	66.7	—	53.6
SDSTrack [14]	CVPR 2024	256 \times 256	—	—	—	—	84.8	62.5	66.5	—	53.1
OneTracker [13]	CVPR 2024	384 \times 384	—	—	—	—	85.7	64.2	67.2	—	53.8
CKD [38]	ACM MM 2024	256 \times 256	93.2	77.2	88.4	65.2	90.0	67.4	73.2	69.3	58.1
PTrMA [43]	TIM 2025	256 \times 256	—	—	86.0	63.4	87.6	65.2	71.5	67.5	56.8
TVTracker [10]	IoTJ 2025	256 \times 256	—	—	87.6	63.7	88.6	64.8	72.6	68.4	57.5
AINet [42]	AAAI 2025	384 \times 384	—	—	87.5	64.8	89.2	67.3	74.2	70.1	59.1
SUTrack [4]	AAAI 2025	384 \times 384	—	—	—	—	92.1	69.2	75.8	—	60.9
STTrack [15]	AAAI 2025	256 \times 256	—	—	—	—	89.8	66.7	76.0	—	60.3
CAFormer [69]	AAAI 2025	256 \times 256	91.8	76.9	85.6	63.2	88.3	66.4	70.0	66.1	55.6
TBSI-Ext [22]	TPAMI 2025	256 \times 256	—	—	—	—	91.0	67.0	75.5	71.5	59.6
CMDTrack [74]	TPAMI 2025	256 \times 256	—	—	—	—	85.9	61.8	68.8	—	56.6
UM-ODTrack [81]	TPAMI 2025	384 \times 384	—	—	—	—	91.5	69.2	74.4	—	58.8
XTrack [57]	ICCV 2025	256 \times 256	—	—	—	—	87.4	64.9	69.1	—	55.7
SMSTracker [3]	ICCV 2025	256 \times 256	—	—	—	—	86.9	64.5	70.3	—	56.0
AETrack [86]	TCSVT 2025	256 \times 256	—	—	90.4	66.3	91.6	68.8	74.7	71.0	59.6
MambaVT [19]	TCSVT 2025	256 \times 256	94.1	75.3	88.0	63.7	88.9	65.8	73.0	69.5	57.9
MoETrack [60]	TIP 2025	256 \times 256	93.6	78.4	—	—	88.1	65.1	72.1	—	57.8
DMD [16]	TIP 2025	256 \times 256	94.2	78.6	89.6	65.8	90.7	67.8	73.6	69.7	58.2
RAGTrack	Ours	256\times256	95.1	79.3	93.2	67.1	93.8	69.5	76.8	73.0	61.1

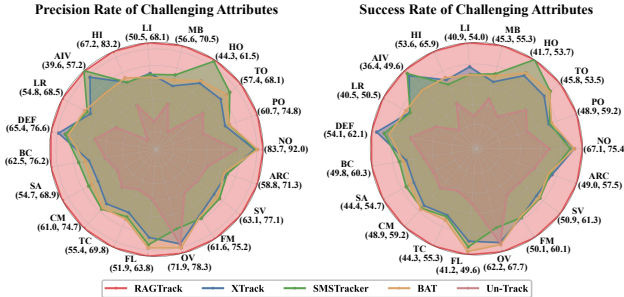


Figure 4. Attribute-based evaluations on the LasHeR dataset.

employs CLIP [48], while the visual backbone uses HiViT-B initialized from SOT [31]. Training employs the AdamW optimizer [36] with a learning rate 10^{-4} and a weight decay 10^{-4} . Template images and search regions are processed at fixed resolutions of 128×128 and 256×256 , respectively. We configure feature dimension $C = 512$, text token $N_h = 1$, and reasoning token $N_r = 1$. Our ATF is deployed at the 6-th, 12-th, 18-th and 24-th layers of HiViT-B with retention ratio $\gamma = 85\%$. We select 256 channels for exchange per modality. In CRM, the local knowledge base maintains $n = 4$ entries, and the retrieval selects $k = 2$ fea-

tures with threshold $\lambda = 1.0$. The MLLM adopts QWen2.5-VL-3B [64]. We train RAGTrack on the LasHeR training set with loss weights $\lambda_{iou} = 2$ and $\lambda_{L_1} = 5$ to balance the regression and classification terms.

4.3. Comparison with State-of-the-Art Trackers

We evaluate RAGTrack on four RGBT benchmarks against state-of-the-art methods. Quantitative results are summarized in Tab. 1, with detailed analysis provided below.

GTOT. Our method delivers state-of-the-art performance with MPR 95.1% in Tab. 1. These results demonstrate significant improvements: +1.5% MPR over MoETrack [60] and +4.0% MSR over MambaVT [19]. The performance gain highlights the effectiveness of our method in handling complex scenarios with modality discrepancies.

RGBT210. Our method shows outstanding performance, achieving 93.2% PR and 67.1% SR in Tab. 1. It exceeds AETrack [86] by a margin of +2.8% PR and shows improved robustness than AINet [42] with a margin of +2.3% SR. These results indicate the robustness of our language-aware representation in improving tracking accuracy across diverse environmental conditions.

Table 2. Ablation study of key components on RGBT234.

Method	MPR \uparrow	MSR \uparrow
Baseline	87.9	64.5
w/ CRM*	89.1	65.0
w/ MTE+CRM*	91.1	66.7
w/ MTE+CRM	91.8	67.4
w/ MTE+CRM+ATF	93.8	69.5

Table 3. Comparison of fusion positions on RGBT234.

6	12	18	24	MPR \uparrow	MSR \uparrow
✓				92.2	67.9
✓	✓			92.5	68.5
✓	✓	✓		92.9	68.8
✓	✓	✓	✓	93.8	69.5

Table 4. Comparison of fusion paradigms on RGBT234.

Module	MPR \uparrow	MSR \uparrow	Params \downarrow
TBSI [17]	92.8	67.6	145.9M
BSI [15]	93.1	68.2	103.6M
DFM [42]	92.7	67.8	110.3M
ATF (Ours)	93.8	69.5	101.8M

RGBT234. Our method achieves 93.8% MPR and 69.5% MSR in Tab. 1, outperforming SMSTracker [3] by 6.9% MPR and STTrack [15] by 2.8% MSR. This improvement validates the superiority of our RAGTrack in leveraging semantic information for accurate target localization.

LasHeR. Our method achieves leading performance with 76.8% PR and 61.1% SR in Tab. 1, surpassing recent trackers: +4.2% PR over TVTracker [10] and +5.4% SR over XTrack [57]. Fig. 4 further reveals the attribute-specific strengths: +10.7% PR in Total Occlusion (TO) and +5.5% SR in Out-of-View (OV). Notably, these attribute-specific advantages underscore the ability of our CRM in maintaining target identities under appearance variations.

4.4. Ablation Studies

We conduct thorough ablation studies to evaluate key components and parameters in RAGTrack, with results summarized in the corresponding tables and figures below.

Component Analysis. We conduct ablation studies to evaluate the contribution of key components. As detailed in Tab. 2, the baseline with a backbone and convolutional fusion achieves 87.9% MPR. The text-free CRM* brings moderate improvement to 89.1% MPR, while incorporating MTE with CRM* further increases performance to 91.1% MPR. The complete CRM yields more clear gains at 91.8% MPR, confirming the value of language information. The full integration of MTE, CRM and ATF achieves superior performance at 93.8% MPR, validating their complementary roles in semantic reasoning and feature interaction. These results indicate that both linguistic enhancement and cross-modal fusion are essential for robust RGBT tracking.

Effect of Fusion Positions. As shown in Tab. 3, fusions at Layers 6, 12, 18 and 24 achieve the optimal perfor-

Table 5. Comparison of augmentation mechanisms on RGBT234.

Module	MPR \uparrow	MSR \uparrow	Params \downarrow
Add	92.5	68.5	101.8M
Transformer	93.3	69.0	104.8M
Mamba	92.7	68.1	103.5M
Ours	93.8	69.5	101.8M

Table 6. Comparison of token configurations on RGBT234.

Reasoning Token			Learnable Token		
N_r	MPR \uparrow	MSR \uparrow	Length	MPR \uparrow	MSR \uparrow
1	93.8	69.5	0	92.8	68.4
2	93.5	69.4	2	93.8	69.5
4	93.2	68.9	4	92.6	68.7

Table 7. Comparison of attention scores on RGBT234.

A_{x2r}	A_{x2h}	A_{x2z}	A_{x2x}	MPR \uparrow	MSR \uparrow
✓				92.0	67.8
✓	✓			92.9	68.7
✓	✓	✓		93.7	69.0
✓	✓	✓	✓	93.8	69.5

mance. Early fusions capture low-level visual features but lack semantics, while deeper fusions extract high-level contextual information but miss fine-grained details. The progressive fusion achieves complementary advantages, showing the importance of cross-layer semantic fusion.

Effect of Different Fusion Paradigms. To validate the effectiveness of ATF, we compare it with several representative fusion paradigms: TBSI [17] establishes template-bridged interaction, BSI [15] introduces template-temporal similarity for background suppression and DFM [42] models inter-modal differences for complementary fusion. In contrast, ATF selects target-relevant tokens and exchanges informative channels. As shown in Tab. 4, ATF achieves the best performance (69.5% MSR) with the fewest trainable parameters (101.8M). The results confirm that ATF enables more discriminative and efficient cross-modal fusion.

Effect of Different Augmentation Mechanisms. We evaluate different designs for the augmentation in CRM. As shown in Tab. 5, our method achieves superior performance. The addition operation yields a limited accuracy due to insufficient representational capacity. Transformer increases trainable parameters by 3.0%, whereas Mamba shows lower performance. These results confirm that our design provides the balance for temporal propagation.

Effect of Reasoning Token Configurations. We evaluate the effect of reasoning token quantity N_r . As shown in Tab. 6, the best result is achieved at $N_r = 1$, outperforming multi-token configurations. This improvement shows that a single reasoning token can capture essential temporal context, while additional tokens may introduce redundant parameters without substantial performance benefits. In addition, our design effectively maintains temporal coherence while minimizing computational overhead.

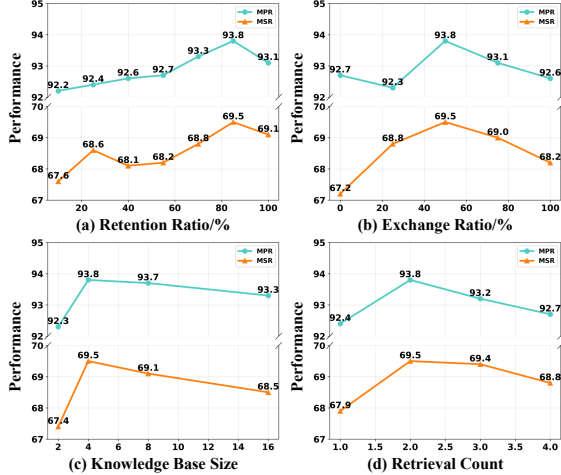


Figure 5. Comparison with different hyper-parameters.

Effect of Learnable Token Lengths. We analyze the effect of learnable token lengths in the sequence prefix. As shown in Tab. 6, the 2-token configuration achieves the optimal performance. The 0-token variant suffers a 1.0% MPR drop, highlighting its limited capacity to sequence-specific semantic variations. The 4-token variant exhibits a 0.8% MSR decline, indicating that excessive tokens introduce redundant information. These results confirm the necessity of contextual tokens for accurate target representation.

Effect of Attention Scores in ATF. We investigate the contribution of different attention scores in ATF. As presented in Tab. 7, performance improves progressively as more attention scores are incorporated. The complete configuration utilizing all four attention scores achieves the optimal results. Our experiments confirm that our multi-score fusion strategy enables effective token selection.

Effect of the Retention Ratio in ATF. The retention ratio determines the proportion of search tokens selected for ATF. As shown in Fig. 5 (a), $\gamma = 85\%$ achieves the leading results with 93.8% MPR. Lower values result in insufficient target information retention, impairing the localization accuracy. In contrast, higher values retain noisy or redundant tokens, degrading feature discrimination. These findings demonstrate that our setting preserves critical target information while filtering noise.

Effect of the Exchange Ratio in ATF. The exchange ratio σ plays a key role in cross-modal interaction. According to Fig. 5 (b), it yields peak performance with $\sigma = 50\%$. At $\sigma = 0$, poor modality fusion results in suboptimal performance. As σ increases to 50%, the improved feature representation enhances target modeling. However, further increasing σ to 100% introduces excessive cross-modal interference, leading to performance degradation. This validates $\sigma = 50\%$ as the optimal configuration, achieving a balance between modality specificity and shared representation.

Effect of the Knowledge Base Size in CRM. We evaluate how the knowledge base size n affects the tracking ac-

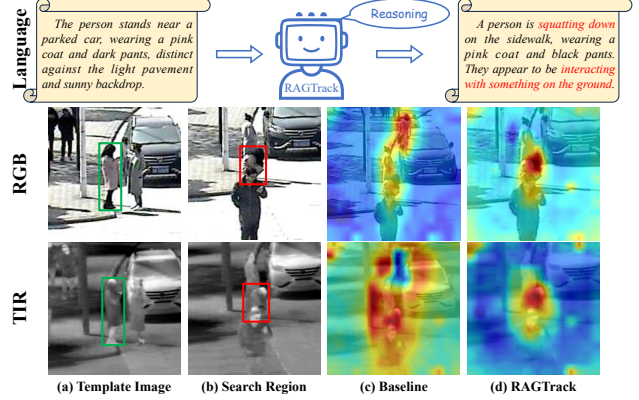


Figure 6. Visualization of attention maps.

curacy. As evidenced by Fig. 5 (c), setting $n = 4$ achieves 69.5% MSR. When $n = 2$, the model lacks sufficient historical context, resulting in a limited reasoning ability. While expanding to $n = 8$ maintains competitive results, further increasing to $n = 16$ causes performance degradation. This decline is likely due to the introduction of redundant or less relevant features. Thus, $n = 4$ strikes the trade-off between context richness and feature quality.

Effect of the Retrieval Count in CRM. As shown in Fig. 5 (d), our method achieves outstanding results with $k = 2$. Smaller counts compromise performance due to limited contextual information, while larger counts perform worse as excessive features introduce noise. By retrieving two relevant features, our method ensures sufficient contextual guidance while achieving effective noise suppression.

4.5. Visualization Analysis

Fig. 6 visualizes attention maps across the search region, illustrating how our RAGTrack improves target localization. The baseline model shows a noticeable discrepancy between the attentive region and the actual target location. In contrast, our proposed method eliminates this misalignment, enhancing the tracking accuracy.

5. Conclusion

In this work, we propose a novel language-aware framework called RAGTrack for RGBT tracking. To establish effective feature representation, we introduce a Multi-modal Transformer Encoder (MTE) that aligns visual and linguistic features through unified modeling. To reduce search redundancies and modality gaps, we design an Adaptive Token Fusion (ATF), which dynamically selects target-relevant tokens and enables adaptive channel exchange. Furthermore, we develop a Context-aware Reasoning Module (CRM) that maintains a dynamic knowledge base and performs context-aware linguistic reasoning to enhance temporal modeling. Extensive experiments on four tracking benchmarks validate the effectiveness of our method.

Acknowledgements

This work was supported partly by the National Natural Science Foundation of China (No. 62441231,62576069), Dalian Science and Technology Innovation Fund(No.2023JJ11CG001) and Natural Science Foundation of Liaoning Province(No.2025-MS-025,2023JH26/10200016).

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [2] Bing Cao, Junliang Guo, Pengfei Zhu, and Qinghua Hu. Bi-directional adapter for multimodal tracking. In *AAAI*, pages 927–935, 2024. 2, 6
- [3] Sixian Chan, Zedong Li, Wenhao Li, Shijian Lu, Chunhua Shen, and Xiaoqin Zhang. Smstracker: Tri-path score mask sigma fusion for multi-modal tracking. In *ICCV*, pages 4766–4775, 2025. 6, 7
- [4] Xin Chen, Ben Kang, Wanting Geng, Jiawen Zhu, Yi Liu, Dong Wang, and Huchuan Lu. Sutrack: Towards simple and unified single object tracking. In *AAAI*, pages 2239–2247, 2025. 6
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 4
- [6] Qi Feng, Vitaly Ablavsky, Qinxun Bai, and Stan Sclaroff. Siamese natural language tracker: Tracking by natural language descriptions with siamese trackers. In *CVPR*, pages 5851–5860, 2021. 1, 2
- [7] Xiaokun Feng, Shiyu Hu, Xuchen Li, Dailing Zhang, Meiqi Wu, Jing Zhang, Xiaotang Chen, and Kaiqi Huang. Atctrack: Aligning target-context cues with dynamic target states for robust vision-language tracking. In *ICCV*, pages 19850–19861, 2025. 2
- [8] Xiaokun Feng, Dailing Zhang, Shiyu Hu, Xuchen Li, Meiqi Wu, Jing Zhang, Xiaotang Chen, and Kaiqi Huang. CSTrack: Enhancing RGB-X tracking via compact spatiotemporal features. In *ICML*, pages 17030–17048, 2025. 2
- [9] Xiaokun Feng, Dailing Zhang, Shiyu Hu, Xuchen Li, Meiqi Wu, Jing Zhang, Xiaotang Chen, and Kaiqi Huang. Enhancing vision-language tracking by effectively converting textual cues into visual cues. In *ICASSP*, pages 1–5, 2025. 2
- [10] Fang Gao, Wenjie Wu, Yan Jin, Jingfeng Tang, Hanbo Zheng, Shengheng Ma, and Jun Yu. Tvtracker: Target-adaptive text-guided visual fusion for multi-modal rgb-t tracking. *IEEE IoTJ*, 12(13):24819–24833, 2025. 1, 6, 7
- [11] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 3
- [12] Mingzhe Guo, Zhipeng Zhang, Liping Jing, Haibin Ling, and Heng Fan. Divert more attention to vision-language object tracking. *IEEE TPAMI*, 46(12):8600–8618, 2024. 3
- [13] 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*, pages 19079–19091, 2024. 2, 6
- [14] 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*, pages 26551–26561, 2024. 2, 6
- [15] 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. In *AAAI*, pages 3581–3589, 2025. 2, 6, 7
- [16] Yufan Hu, Zekai Shao, Bin Fan, and Hongmin Liu. Dual-level modality de-biasing for RGBT tracking. *IEEE TIP*, 34: 2667–2679, 2025. 6
- [17] Tianrui Hui, Zizheng Xun, Fengguang Peng, Junshi Huang, Xiaoming Wei, Xiaolin Wei, Jiao Dai, Jizhong Han, and Si Liu. Bridging search region interaction with template for RGBT tracking. In *CVPR*, pages 13630–13639, 2023. 2, 6, 7
- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456, 2015. 5
- [19] Simiao Lai, Chang Liu, Jiawen Zhu, Ben Kang, Yang Liu, Dong Wang, and Huchuan Lu. MambaVT: Spatio-temporal contextual modeling for robust RGBT tracking. *IEEE TCSVT*, 35(9):9312–9323, 2025. 6
- [20] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *ECCV*, pages 734–750, 2018. 5
- [21] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. In *NeurIPS*, pages 9459–9474, 2020. 3
- [22] Bo Li, Fengguang Peng, Tianrui Hui, Xiaoming Wei, Xiaolin Wei, Lijun Zhang, Hang Shi, and Si Liu. Rgb-t tracking with template-bridged search interaction and target-preserved template updating. *IEEE TPAMI*, 47(1):634–649, 2025. 6
- [23] Chenglong Li, Hui Cheng, Shiyi Hu, Xiaobai Liu, Jin Tang, and Liang Lin. Learning collaborative sparse representation for grayscale-thermal tracking. *IEEE TIP*, 25(12):5743–5756, 2016. 5
- [24] Chenglong Li, Nan Zhao, Yijuan Lu, Chengli Zhu, and Jin Tang. Weighted sparse representation regularized graph learning for RGBT object tracking. In *ACM MM*, pages 1856–1864, 2017. 5
- [25] Chenglong Li, Xinyan Liang, Yijuan Lu, Nan Zhao, and Jin Tang. Rgbt object tracking: Benchmark and baseline. *PR*, 96:106977, 2019. 5
- [26] Chenglong Li, Lei Liu, Andong Lu, Qing Ji, and Jin Tang. Challenge-aware rgbt tracking. In *ECCV*, pages 222–237, 2020. 6

- [27] 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. *IEEE TIP*, 31:392–404, 2021. 5
- [28] Hao Li, Yuhao Wang, Xiantao Hu, Wenning Hao, Pingping Zhang, Dong Wang, and Huchuan Lu. Cadtrack: Learning contextual aggregation with deformable alignment for robust rgbt tracking. *arXiv preprint arXiv:2511.17967*, 2025. 1
- [29] Xin Li, Yuqing Huang, Zhenyu He, Yaowei Wang, Huchuan Lu, and Ming-Hsuan Yang. Citetracker: Correlating image and text for visual tracking. In *ICCV*, pages 9974–9983, 2023. 1, 2
- [30] Xuchen Li, Xiaokun Feng, Shiyu Hu, Meiqi Wu, Dailing Zhang, Jing Zhang, and Kaiqi Huang. Dtlm-vlt: Diverse text generation for visual language tracking based on llm. In *CVPRW*, pages 7283–7292, 2024. 3
- [31] Xiaohai Li, Bineng Zhong, Qihua Liang, Zhiyi Mo, Jian Nong, and Shuxiang Song. Dynamic updates for language adaptation in visual-language tracking. In *CVPR*, pages 19165–19174, 2025. 3, 6
- [32] Zhenyang Li, Ran Tao, Efstratios Gavves, Cees GM Snoek, and Arnold WM Smeulders. Tracking by natural language specification. In *CVPR*, pages 6495–6503, 2017. 2
- [33] Bingxi Liu, Calvin Chen, Junhao Li, Guyang Yu, Haoqian Song, Xuchen Liu, Jinqiang Cui, and Hong Zhang. Trackingmim: Efficient mamba-in-mamba serialization for real-time uav object tracking. *arXiv preprint arXiv:2507.01535*, 2025. 3
- [34] Chang Liu, Xiao-Fan Chen, Chun-Juan Bo, and Dong Wang. Long-term visual tracking: review and experimental comparison. *MIR*, 19(6):512–530, 2022.
- [35] Chang Liu, Yongsheng Yuan, Xin Chen, Huchuan Lu, and Dong Wang. Spatial-temporal initialization dilemma: towards realistic visual tracking. *VI*, 2(1):35, 2024. 1
- [36] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 6
- [37] Andong Lu, Chenglong Li, Yuqing Yan, Jin Tang, and Bin Luo. Rgbt tracking via multi-adapter network with hierarchical divergence loss. *IEEE TIP*, 30:5613–5625, 2021. 2
- [38] Andong Lu, Jiacong Zhao, Chenglong Li, Yun Xiao, and Bin Luo. Breaking modality gap in RGBT tracking: Coupled knowledge distillation. In *ACM MM*, pages 9291–9300, 2024. 1, 6
- [39] Andong Lu, Chenglong Li, Jiacong Zhao, Jin Tang, and Bin Luo. Modality-missing rgbt tracking: Invertible prompt learning and high-quality benchmarks. *IJCV*, 133(5):2599–2619, 2025.
- [40] Andong Lu, Cun Qian, Chenglong Li, Jin Tang, and Liang Wang. Duality-gated mutual condition network for rgbt tracking. *IEEE TNLS*, 36(3):4118–4131, 2025.
- [41] Andong Lu, Wanyu Wang, Chenglong Li, Jin Tang, and Bin Luo. After: Attention-based fusion router for rgbt tracking. *IEEE TIP*, 34:4386–4401, 2025.
- [42] Andong Lu, Wanyu Wang, Chenglong Li, Jin Tang, and Bin Luo. RGBT tracking via all-layer multimodal interactions with progressive fusion Mamba. In *AAAI*, pages 5793–5801, 2025. 2, 6, 7
- [43] Binxin Luo, Dongxu Liu, Xianrong Peng, Haorui Zuo, Jianlin Zhang, Meihui Li, and Yuxing Wei. Progressive transformer with multi-modality adaptation for rgb-t tracking. *IEEE TIM*, 74:1–17, 2025. 1, 6
- [44] Yinchao Ma, Yuyang Tang, Wenfei Yang, Tianzhu Zhang, Jinpeng Zhang, and Mengxue Kang. Unifying visual and vision-language tracking via contrastive learning. In *AAAI*, pages 4107–4116, 2024. 3
- [45] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, pages 807–814, 2010. 5
- [46] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, pages 4293–4302, 2016. 2
- [47] Haolin Qin, Tingfa Xu, Tianhao Li, Zhenxiang Chen, Tao Feng, and Jianan Li. Must: The first dataset and unified framework for multispectral uav single object tracking. In *CVPR*, pages 16882–16891, 2025. 1
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 6
- [49] Hamid Rezaatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *CVPR*, pages 658–666, 2019. 5
- [50] Zekai Shao, Yufan Hu, Bin Fan, and Hongmin Liu. PURA: Parameter update-recovery test-time adaption for RGBT tracking. In *CVPR*, pages 22089–22098, 2025. 1
- [51] Zikai Song, Junqing Yu, Yi-Ping Phoebe Chen, and Wei Yang. Transformer tracking with cyclic shifting window attention. In *CVPR*, pages 8791–8800, 2022.
- [52] Zikai Song, Run Luo, Junqing Yu, Yi-Ping Phoebe Chen, and Wei Yang. Compact transformer tracker with correlative masked modeling. In *AAAI*, pages 2321–2329, 2023.
- [53] Zikai Song, Ying Tang, Run Luo, Lintao Ma, Junqing Yu, Yi-Ping Phoebe Chen, and Wei Yang. Autogenic language embedding for coherent point tracking. In *ACM MM*, pages 2021–2030, 2024.
- [54] Zikai Song, Run Luo, Lintao Ma, Ying Tang, Yi-Ping Phoebe Chen, Junqing Yu, and Wei Yang. Temporal coherent object flow for multi-object tracking. In *AAAI*, pages 6978–6986, 2025.
- [55] Yiming Sun, Fan Yu, Shaoxiang Chen, Yu Zhang, Junwei Huang, Yang Li, Chenhui Li, and Changbo Wang. Chat-tracker: Enhancing visual tracking performance via chatting with multimodal large language model. In *NeurIPS*, pages 39303–39324, 2024. 3
- [56] Yuedong Tan, Jiawei Shao, Eduard Zamfir, Ruanjun Li, Zhaochong An, Chao Ma, Danda Pani Paudel, Luc Van Gool, Radu Timofte, and Zongwei Wu. What you have is what you track: Adaptive and robust multimodal tracking. In *ICCV*, pages 3455–3465, 2025. 2
- [57] Yuedong Tan, Zongwei Wu, Yuqian Fu, Zhuyun Zhou, Guolei Sun, Eduard Zamfi, Chao Ma, Danda Pani Paudel, Luc Van Gool, and Radu Timofte. XTrack: Multimodal

- training boosts RGB-X video object trackers. In *ICCV*, pages 5734–5744, 2025. 1, 2, 6, 7
- [58] Xiaqiang Tang, Qiang Gao, Jian Li, Nan Du, Qi Li, and Sihong Xie. MBA-RAG: a bandit approach for adaptive retrieval-augmented generation through question complexity. In *COLING*, pages 3248–3254, 2025. 3
- [59] Zhangyong Tang, Tianyang Xu, Xiaojun Wu, Xue-Feng Zhu, and Josef Kittler. Generative-based fusion mechanism for multi-modal tracking. In *AAAI*, pages 5189–5197, 2024. 6
- [60] Zhangyong Tang, Tianyang Xu, Xiao-Jun Wu, Xuefeng Zhu, Chunyang Cheng, Zhenhua Feng, and Josef Kittler. Revisiting rgbt tracking benchmarks from the perspective of modality validity: A new benchmark, problem, and solution. *IEEE TIP*, 34:7235–7249, 2025. 6
- [61] 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*, pages 7064–7073, 2020. 2, 6
- [62] Hongyu Wang, Xiaotao Liu, Yifan Li, Meng Sun, Dian Yuan, and Jing Liu. Temporal adaptive rgbt tracking with modality prompt. In *AAAI*, pages 5436–5444, 2024. 2, 6
- [63] Mengmeng Wang, Teli Ma, Shuo Xin, Xiaojun Hou, Jiazheng Xing, Guang Dai, Jingdong Wang, and Yong Liu. Visual object tracking across diverse data modalities: A review. *arXiv preprint arXiv:2412.09991*, 2024. 2
- [64] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 6
- [65] Xiao Wang, Xiujun Shu, Zhipeng Zhang, Bo Jiang, Yaowei Wang, Yonghong Tian, and Feng Wu. Towards more flexible and accurate object tracking with natural language: Algorithms and benchmark. In *CVPR*, pages 13763–13773, 2021. 2
- [66] Xiao Wang, Liye Jin, Xufeng Lou, Shiao Wang, Lan Chen, Bo Jiang, and Zhipeng Zhang. Reasoningtrack: Chain-of-thought reasoning for long-term vision-language tracking. *arXiv preprint arXiv:2508.05221*, 2025. 3
- [67] 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*, pages 19156–19166, 2024. 6
- [68] Yun Xiao, Mengmeng Yang, Chenglong Li, Lei Liu, and Jin Tang. Attribute-based progressive fusion network for rgbt tracking. In *AAAI*, pages 2831–2838, 2022. 6
- [69] Yun Xiao, Jiacong Zhao, Andong Lu, Chenglong Li, Bing Yin, Yin Lin, and Cong Liu. Cross-modulated attention transformer for RGBT tracking. In *AAAI*, pages 8682–8690, 2025. 6
- [70] 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*, pages 341–357, 2022. 4
- [71] Chunhui Zhang, Xin Sun, Yiqian Yang, Li Liu, Qiong Liu, Xi Zhou, and Yanfeng Wang. All in one: Exploring unified vision-language tracking with multi-modal alignment. In *ACM MM*, pages 5552–5561, 2023. 2
- [72] Huanlong Zhang, Jingchao Wang, Jianwei Zhang, Tianzhu Zhang, and Bineng Zhong. One-stream vision-language memory network for object tracking. *IEEE TMM*, 26:1720–1730, 2023. 2
- [73] Tianlu Zhang, Hongyuan Guo, Qiang Jiao, Qiang Zhang, and Jungong Han. Efficient rgb-t tracking via cross-modality distillation. In *CVPR*, pages 5404–5413, 2023. 6
- [74] Tianlu Zhang, Qiang Zhang, Kurt Debattista, and Jungong Han. Cross-modality distillation for multi-modal tracking. *IEEE TPAMI*, 47(7):5847–5865, 2025. 6
- [75] Xiaosong Zhang, Yunjie Tian, Lingxi Xie, Wei Huang, Qi Dai, Qixiang Ye, and Qi Tian. Hivit: A simpler and more efficient design of hierarchical vision transformer. In *ICLR*, 2023. 3
- [76] Jie Zhao, Jingshu Zhang, Dongdong Li, and Dong Wang. Vision-based anti-uav detection and tracking. *IEEE TITS*, 23(12):25323–25334, 2022. 1
- [77] Yaozong Zheng, Bineng Zhong, Qihua Liang, Zhenjun Tang, Rongrong Ji, and Xianxian Li. Leveraging local and global cues for visual tracking via parallel interaction network. *IEEE TCSVT*, 33(4):1671–1683, 2022.
- [78] Yaozong Zheng, Bineng Zhong, Qihua Liang, Guorong Li, Rongrong Ji, and Xianxian Li. Toward unified token learning for vision-language tracking. *IEEE TCSVT*, 34(4):2125–2135, 2023. 2
- [79] Yaozong Zheng, Bineng Zhong, Qihua Liang, Zhiyi Mo, Shengping Zhang, and Xianxian Li. Odtrack: Online dense temporal token learning for visual tracking. In *AAAI*, pages 7588–7596, 2024.
- [80] Yaozong Zheng, Bineng Zhong, Qihua Liang, Ning Li, and Shuxiang Song. Decoupled spatio-temporal consistency learning for self-supervised tracking. In *AAAI*, pages 10635–10643, 2025.
- [81] Yaozong Zheng, Bineng Zhong, Qihua Liang, Shengping Zhang, Guorong Li, Xianxian Li, and Rongrong Ji. Towards universal modal tracking with online dense temporal token learning. *IEEE TPAMI*, 47(11):10192–10209, 2025. 1, 6
- [82] Li Zhou, Zikun Zhou, Kaige Mao, and Zhenyu He. Joint visual grounding and tracking with natural language specification. In *CVPR*, pages 23151–23160, 2023. 2
- [83] Jiawen Zhu, Simiao Lai, Xin Chen, Dong Wang, and Huchuan Lu. Visual prompt multi-modal tracking. In *CVPR*, pages 9516–9526, 2023. 2, 6
- [84] Xue-Feng Zhu, Tianyang Xu, Yifan Pan, Jinjie Gu, Xi Li, Jiwen Lu, Xiao-Jun Wu, and Josef Kittler. Collaborating vision, depth, and thermal signals for multi-modal tracking: Dataset and algorithm. In *NeurIPS*, 2025. 2
- [85] Yabin Zhu, Chenglong Li, Bin Luo, Jin Tang, and Xiao Wang. Dense feature aggregation and pruning for rgbt tracking. In *ACM MM*, pages 465–472, 2019. 2
- [86] Zhiruo Zhu, Bineng Zhong, Qihua Liang, Hongtao Yang, Yaozong Zheng, and Ning Li. Adaptive expert decision for rgb-t tracking. *IEEE TCSVT*, 35(10):10330–10338, 2025. 6