

# Single-Model and Any-Modality for Video Object Tracking

Zongwei Wu<sup>1</sup> Jilai Zheng<sup>2</sup> Xiangxuan Ren<sup>2</sup> Florin-Alexandru Vasluianu<sup>1</sup>  
 Chao Ma<sup>2\*</sup> Danda Pani Paudel<sup>3</sup> Luc Van Gool<sup>3,4</sup> Radu Timofte<sup>1</sup>

<sup>1</sup> Computer Vision Lab, CAIDAS & IFI, University of Wurzburg <sup>2</sup> AI Institute, Shanghai Jiao Tong University  
<sup>3</sup> INSAIT, Sofia University <sup>4</sup> CVL, ETH Zurich

## Abstract

In the realm of video object tracking, auxiliary modalities such as depth, thermal, or event data have emerged as valuable assets to complement the RGB trackers. In practice, most existing RGB trackers learn a single set of parameters to use them across datasets and applications. However, a similar single-model unification for multi-modality tracking presents several challenges. These challenges stem from the inherent heterogeneity of inputs – each with modality-specific representations, the scarcity of multi-modal datasets, and the absence of all the modalities at all times. In this work, we introduce *Un-Track*, a *Unified Tracker* of a single set of parameters for any modality. To handle any modality, our method learns their common latent space through low-rank factorization and reconstruction techniques. More importantly, we use only the RGB-X pairs to learn the common latent space. This unique shared representation seamlessly binds all modalities together, enabling effective unification and accommodating any missing modality, all within a single transformer-based architecture. Our *Un-Track* achieves **+8.1 absolute F-score gain**, on the *DepthTrack* dataset, by introducing only +2.14 (over 21.50) GFLOPs with +6.6M (over 93M) parameters, through a simple yet efficient prompting strategy. Extensive comparisons on five benchmark datasets with different modalities show that *Un-Track* surpasses both SOTA unified trackers and modality-specific counterparts, validating our effectiveness and practicality. The source code is publicly available at <https://github.com/Zongwei97/UnTrack>.

## 1. Introduction

Video object tracking [9, 12, 86, 90] is a fundamental task in computer vision with wide-ranging applications spanning from surveillance [89] to augmented reality [23], where accuracy and robustness are paramount. While traditional

\* Corresponding author.

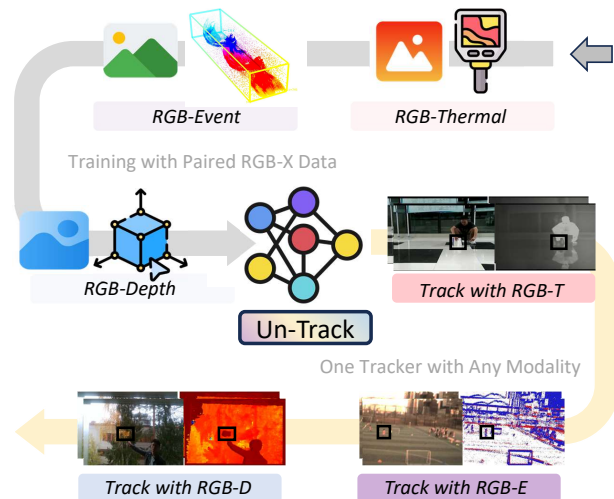


Figure 1. Un-Track is a unified tracker with a single parameter set that seamlessly integrates any modality (of RGB-X).

RGB trackers have shown promise in general settings, they often struggle with challenging scenarios such as occlusions [40, 47], low visibility [83, 92], or fast-moving objects [10, 48]. For more reliable tracking under such challenging conditions, the integration of auxiliary modalities (X) like depth [20, 68], thermal [83, 87], and event [79, 98] have proven effective in multimodal tracking.

While the idea of fusing RGB with other modalities holds promise [82, 84], the main challenge is the discrepancy in the representation of information across different modalities. Despite the success of previous fusion works [15, 60] tailored for each specific scenario to improve RGB trackers, their reliance on modality-specific designs limits adaptability. Recent initiatives [69, 91] towards a uniform architecture for various modalities show promise but necessitate modality-specific fine-tuning. This approach leads to multiple parameter sets, as shown in Fig. 1(a), thereby compromising practicality in diverse real-world applications.

In this work, we aim to avoid such modality-specific fine-tuning to keep only one model-parameter set at all times. Another practical constraint arises from the differ-

ences in available auxiliary modalities across settings. Unifying modalities by a common representation can handle any modality at its disposal, addressing the mentioned problems by its virtue. However, two additional multifaceted challenges emerge from the scarcity of multimodal datasets and the absence of all paired data combinations. The former makes cross-modal mappings through a large-scale data prior unfeasible [19, 39, 76], while the latter leads to missing modalities and renders joint learning using all possible combinations of paired data unfeasible [41, 80].

To achieve such a unification, in this paper, we present **Un-Track**, denoting “one” in French, which learns a cohesive embedding across diverse input modalities. Unlike previous approaches [51, 58, 73, 80], Un-Track relies solely on RGB-X pairs for training, with X representing auxiliary modalities, without the need for all modalities to co-occur. Our objective is to discover a shared embedding seamlessly binding all auxiliary modalities (as depicted in Fig. 1(b)). More specifically, we leverage the factorization prior, allowing reasoning about a common embedding directly from the low-rank latent space. Factorization is a simple composition prior with the assumption that the global approximation can be constructed from a set of subset vectors. The factorization prior, successfully utilized in previous studies [1, 4, 24], is employed in our work to reconstruct a shared embedding. This process transforms the heterogeneous modal representation into a uniform one, thereby facilitating the emergent cross-modal alignment.

Moreover, to harness the full potential of auxiliary inputs while maintaining efficiency, we leverage cross-modal features as prompts to enable RGB-X interaction. Different from previous works [25, 27], our goal is to enhance less reliable tokens, defined by a learnable score function, using multimodal cues. We approach this as a token recovery problem and leverage low-rank factorization to achieve the goal, which is first suggested in multimodal fusion, to the best of our knowledge. With its unified model architecture and prompting blocks, Un-Track is the first to offer support for cross-modality alignment under a single architecture with uniform parameters. In comparison to our RGB baseline with 21.50 GFLOPs and 92M parameters, Un-Track introduces only +2.14 GFLOPs with +6.6M parameters, resulting in a significant +8.1 absolute F-score gain demonstrated on the DepthTrack dataset. Extensive comparisons across five datasets with different modalities validate Un-Track’s superiority over specialized SOTA models, surpassing both unified trackers and modality-specific fine-tuned counterparts by a substantial margin.

## 2. Related Works

**Multimodal tracking:** Video object tracking [7, 71] aims to detect objects in a video sequence based on their initial positions. Early approaches treated tracking as a per-frame

target-matching problem, with Siamese networks [32, 33, 50, 61, 74, 85] being a notable example. More recently, transformer-based methods [2, 5, 6, 11, 54, 63] have gained popularity for feature extraction and per-frame correlation in tracking. Large-scale training datasets [14, 22, 42, 45] have empowered RGB trackers to uniformly apply parameters across various applications. While RGB trackers deliver promising results, challenges such as occlusion, low illumination, and fast-moving scenes have led to the exploration of additional modalities. Several works have investigated the role of depth [70, 95], thermal [37, 38], and event modalities [49, 55] in enhancing tracking performance. Specifically, depth cues [16, 67] contribute to handling objects with different camera distances; thermal cameras [56, 60] address challenges such as low illumination; event cameras improve the temporal awareness [57, 78, 97].

Despite the plausible achievements, many rely on modality-specific blocks designed for individual modalities [88, 93], limiting their adaptability. Recent efforts have focused on achieving architectural unification [69, 91], yet they still necessitate modality-specific fine-tuning, resulting in distinct parameter sets for different modalities. The ideal scenario would involve a large-scale dataset encompassing all possible modal combinations, but current tracking datasets predominantly feature a single modality — depth [66, 94], thermal [35, 36], or event [55] — posing challenges for a unified model with a single parameter set.

**Learning with Missing modalities:** Recent research has addressed real-world scenarios where models must cope with missing modalities [41, 44]. One common strategy involves estimating missing values by learning joint multimodal representations [31, 75], feasible when complete samples are available during training. However, tracking datasets typically exhibit only one modality at a time, complicating the learning of such joint representations. Other works [19, 39] implicitly learn cross-modal alignment end-to-end using large-scale datasets and deep networks, demanding substantial computational resources. Extending such approaches to tracking is challenging due to limited downstream datasets and real-world applicability constraints [26, 64]. In contrast to existing methods, our approach investigates cross-modal relationships by leveraging edge priors to learn a joint representation that unifies all modalities. Our method does not require the simultaneous occurrence of all modalities, offering a unique perspective on dealing with diverse and individual modalities.

## 3. Methods

### 3.1. Overall Framework

In this paper, our primary focus is on multimodal tracking, with the constraint that only one modality is available at a time, as shown in Fig. 2. We define our multimodal

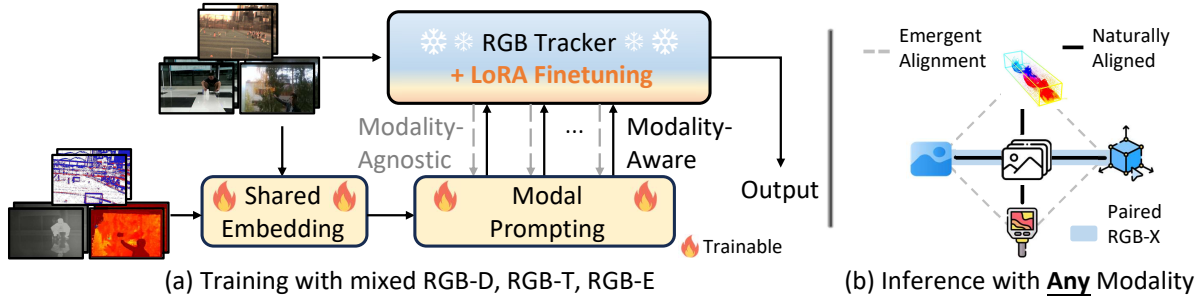


Figure 2. Our **proposed framework**, termed Un-Track, is composed of a shared embedding, a modal prompting, and a LoRA-finetuned pretrained RGB tracker. The shared embedding learns a joint representation that unifies all modalities (Sec. 3.2). The modal prompting block enhances feature modeling with modal awareness at each scale (Sec. 3.3). To track the target object, we finetune a pretrained foundation model [72] using the LoRA technique (Sec. 3.4). Our model achieves a unified model applicable across different modalities under a single parameter set. During inference, Un-Track seamlessly integrates any image-paired data, thanks to the emergent alignment.

tracking dataset as  $M = \{M^D, M^T, M^E\}$ , where  $M^X$  represents the subset dataset with only a single modality  $X$  available. Conventional methods tackling missing modalities often use dummy inputs for the absent pixels [31, 53], simulating complete datasets with all possible combinations of paired datasets. In contrast, our method transforms any auxiliary input into a shared embedding, seamlessly binding all modalities together and creating a complete and paired representation with the master RGB feature.

To mitigate overfitting on sparse downstream multimodal datasets, we adopt a transformer-based RGB tracker with frozen parameters and fine-tune it for multimodal tracking. Leveraging a lightweight outer prompting method, we identify uncertain tokens at each scale and enhance them with cross-modal awareness. Simultaneously, an inner fine-tuning process is implemented using the LoRA technique [21].

During training, our model learns the shared embedding from samples in the mixed dataset  $M$ , effectively binding all modalities together. As for inference, our model can accommodate any modal input  $X$ , thanks to the emergent alignment. Our trainable parameters only include cross-modal binding, outer prompting parameters, and inner LoRA parameters, ensuring a training-friendly pipeline that can be efficiently employed end-to-end on a single 24G GPU.

### 3.2. Shared Embedding

**Explicit Edge Awareness:** We observe that, as illustrated in Fig. 2(a), depth data introduce 3D distance information, effectively delineating objects with varying granularity and enhancing the sharpness of 3D boundaries; thermal images generate a scene heat map, highlighting objects based on their temperatures and providing clearer contours; event data capture intensity changes, particularly around an object’s outbound region. Notably, a consistent feature emerges across these modalities: the representation of the “true” objects’ shape, often manifested through edges.

Motivated by this observation, our objective is to harness

edge embedding to unify all modalities. To achieve this, as shown in Fig. 3, we generate gradient maps from auxiliary modalities by computing differences between neighboring pixels along both the x- and y-axes. Simultaneously, without compromising texture edge, we also generate RGB gradient maps. Subsequently, all gradient maps are integrated with the visual feature, forming the gradient feature  $G$ .

**Implicit Low-Rank Reconstruction:** While edges present a shared feature across different modalities, exclusively transforming all modalities into edges may risk overlooking modality-specific clues. Therefore, we introduce an implicit learning pipeline to complement this by discovering the shared embedding, guided by the previously generated explicit edge awareness. This combined approach allows for a more effective identification of the shared embedding, leveraging both data-driven learning and edge priors.

In technical terms, we redefine the challenge of learning the shared embedding as a quest for the shared low-rank approximation. Both objectives share the essence of distilling common features across all modalities. However, direct estimation of the low-rank approximation becomes impractical due to the distinct data domains and modal representations. In response, we propose a pragmatic strategy: decomposing the shared low-rank vector into the low-rank of each subset component. This alternative, more manageable and feasible within a single domain, lays the groundwork for approximating the global shared low-rank from these individual low-rank components.

The overall algorithm can be found in Algorithm 1. Specifically, let  $M$  be the input feature with mixed auxiliary modalities, decomposed into subset features  $D, T, E$  representing depth, thermal, and event samples from subset datasets  $M^D, M^T, M^E$ . Their respective  $k_{th}$  low-rank matrices  $D_k, T_k, E_k$ , are approximated by:

$$D_k = \sigma_d(D), \quad T_k = \sigma_t(T), \quad E_k = \sigma_e(E), \quad (1)$$

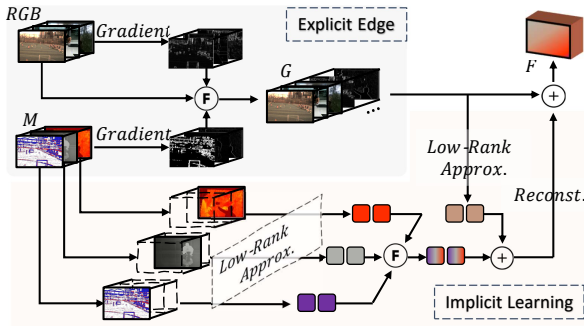
where  $\sigma_x$  is the modality-specific learning through a simple MLP projecting features from input channel  $c$  to the low-

### Algorithm 1 Implicit Shared Embedding

**Input:**  $M$  - Mix Modal Feature,  $G$  - Explicit Gradient Binding.

**Output:**  $F$  - Reconstructed feature

1. **Initialize:** Separate  $M$  into subsets, each representing a modality-specific feature:  $D, T, E$ .
2. **In-Domain Approximation:** Train individual approximators  $\sigma_x$  for each modality  $x$  to derive domain-specific low-rank matrices  $D_k, T_k, E_k$ . Idem for explicit low-rank gradient  $G_k$ .
3. **Fuse and Guide:** Merge  $D_k, T_k, E_k$  through a fusion function  $\phi_1$ . Incorporate the explicit gradient  $G_k$  using addition, after projecting it into a compatible latent space through  $\phi_2$ .
4. **Reconstruct:** Train a reconstruction function  $\Phi_R$  to construct the shared embedding  $F$ , considering the explicit guidance  $G$  in the original feature space.



RGB: Image M: Mixed Modalities G: Gradient Feature F: Shared Embedding

Figure 3. **Shared Embedding.** We derive a joint representation through low-rank factorization and reconstruction. Such an implicit learning is additionally integrated with explicit edge awareness to enhance the embedding.

rank space  $k$  ( $k < c$ ). Simultaneously, we compute the low-rank matrix  $G_k$  from the gradient feature  $G$ .

The global shared low-rank matrix  $M_k$  is then approximated by fusing the subset low-rank matrices  $D_k, T_k, E_k$ , along with the gradient guidance. Technically, we concatenate  $D_k, T_k, E_k$  and learn the joint low-rank approximation, incorporating it with the gradient guidance. This pipeline can be expressed as follows:

$$M_k = \phi_{R_1}([D_k, T_k, E_k]) + \phi_{R_2}(G_k), \quad (2)$$

where  $[\cdot]$  is the channel concatenation and  $\phi_{R_i}$  are the MLP projections to the low-rank latent space. Finally, we reconstruct the shared embedding  $F$  through:

$$F = \Phi_R(M_k) + G, \quad (3)$$

where  $\Phi_R$  is another MLP that projects the jointly-learned low-rank matrix back to the departing feature space. Our ablation studies validate that this subset low-rank approximation and regrouping efficiently unify all input modalities, despite their heterogeneous representations.




Figure 4. **Modal Prompting.** For the visual feature  $I$ , we employ a score function to categorize tokens into negative, uncertain, and positive segments. Using a token exchange policy, we discard negative tokens, enhance uncertain ones with corresponding tokens from  $F$ , and retain positive ones. Then, we transform the feature fusion task into a token recovery problem, addressed by low-rank factorization. Similarly, we extract the most informative low-rank matrix from  $F$  to fuse and reconstruct the shared output.

### 3.3. Outer Modal Prompting

RGB-tracker may fail to perform accurately in corner cases where auxiliary clues can contribute. Drawing inspiration from the success of adapting large pre-trained models to specific downstream tasks [21, 25], we introduce a modal prompting method devised to enhance RGB token  $I$  with modality-awareness  $F$ , as shown in Fig. 4.

Specifically, our approach employs a shrinkage token fusion strategy. Taking  $I$  as an example, we categorize tokens into three groups—negative, uncertain, and positive—based on a dynamic scoring function  $s$ . These regions are defined with mask form, expressed as  $m_n, m_u, m_p = s(I)$ . To harness multimodal clues effectively, we replace negative tokens with those from the other modality, omit uncertain ones with dummy values, and retain the positive tokens. Subsequently, these modified tokens undergo projection into the low-rank space using the approximation function  $\sigma_c$ . Our objective is to enhance robustness by completing uncertain tokens with information from reliable neighboring tokens. Mathematically, we obtain the first low-rank matrix  $I_{l_1}$  by:

$$I_{l_1} = \sigma_c(m_n \cdot F + m_p \cdot I). \quad (4)$$

Next, we target the uncertain tokens by merging them with the paired tokens from the other modality and approximate the low-rank matrix similarly. Here, we aim to throw out the possible noise, resulting in a matrix that is more informative than the original. Let  $\sigma_n$  be another approximation function, we obtain the second low-rank matrix  $I_{l_n}$  by:

$$I_{l_2} = \sigma_n(m_u \cdot F + m_u \cdot I). \quad (5)$$

Then, we fuse these two low-rank matrices in low-rank

	Depth-Specific Parameters							Uni-model with a Single Set of Parameters						
	ATCAIS [28]	DDiMP [28]	DeT [66]	SPT [94]	ProTrack [69]	ViPT [91]	Un-Track (ours)	Stark [62]	AiATrack [17]	OSTrack [72]	UniNext [65]	SeqTrack [8]	ViPT [91]	Un-Track (ours)
F-score(↑)	0.476	0.485	0.532	0.538	0.578	<b>0.594</b>	<b>0.612</b>	0.397	0.515	0.569	0.422	0.590	0.561	<b>0.610</b>
Re(↑)	0.455	0.469	0.506	0.549	0.573	<b>0.596</b>	<b>0.610</b>	0.406	0.526	0.582	0.432	0.600	0.562	<b>0.610</b>
Pr(↑)	0.500	0.503	0.560	0.527	0.583	<b>0.592</b>	<b>0.613</b>	0.388	0.505	0.557	0.413	0.580	0.560	<b>0.610</b>

Table 1. Overall performance on DepthTrack test set [66]. Red/Green/Blue stands for the best/second/third performance. Our depth-specific model sets the SOTA record, while our uni-model with single parameters set also outperforms previous depth-specific SOTA.

space, forming the shared low-rank matrix  $I_l$  for input  $I$ . In such a manner, we improve the image token modeling by fully benefiting from the auxiliary clues. Mathematically, the whole process can be formulated as:

$$I_l = \varphi_P([I_1, I_2]), \quad (6)$$

where  $\varphi_P$  is learnable fusion. Similarly, from the input  $F$ , we follow the same process to obtain the low-rank matrix  $F_l$ . For the cross-modal fusion, we add these two low-rank matrices and then project back to the original space. We can obtain the fused output  $O$  by:

$$O = \Phi_P(I_l + F_l), \quad (7)$$

where  $\Phi_P$  is another learnable MLP.

Our method can be regarded as a mixer of token exchange (for negative tokens) and token fusion (for uncertain tokens), while retaining the most informative modality-specific clues (for positive tokens). As the majority of fusion operations occur in the low-rank feature space, our progressive cross-modal shrinkage avoids imposing a significant additional computational burden, while being able to excavate and accumulate the rich cues from each modality for effective modal prompting.

### 3.4. Inner Finetuning

In addition to the outer modal prompting, we incorporate the LoRA technique [21] for more efficient finetuning. For each transformer attention module with the weight matrix  $W_0 \in \mathbb{R}^{d \times k}$ , we introduce two learnable matrices:  $B \in \mathbb{R}^{d \times r}$  and  $A \in \mathbb{R}^{r \times k}$ . This leads to the replacement of the frozen attention mechanism  $h = W_0 \mathbf{x}$  with the new LoRA attention:

$$h = W_0 \mathbf{x} + B \mathbf{A} \mathbf{x}. \quad (8)$$

To train our network, we adopt the same loss functions as our baseline tracker [72] for end-to-end learning.

## 4. Experiments

### 4.1. Training Data

In the absence of a comprehensive multi-modal tracking dataset encompassing all possible combinations (RGB-D-T-E), we have only one RGB+X (i.e. RGB+depth, **or**

RGB+thermal, **or** RGB+event) at a time. The conventional modality-specific adapted settings are referred as “X-Specific”, whereas the main target of this paper with one model trained on all modality pairs is called “Uni-model” (or a single set of parameters). The training and evaluation settings are summarized as follows:

Model	Trained on	Evaluated on
X-Specific	$\text{RGB}_i + X_i$	$\text{RGB}_i + X_i$
Uni-model	$\bigcup_i \text{RGB}_i + X_i$	$\text{RGB}_i + X_i$

Our RGB-D samples are sourced from DepthTrack [66], a pioneering RGB-D tracking benchmark with 150 training long-term sequences. RGB-T samples are extracted from the extensive LasHeR [36] dataset, featuring 979 diverse training sequences. RGB-E samples are obtained from VisEvent [55], which boasts 500 real-world sequences. Each D/T/E input is transformed into an RGB-like form.

### 4.2. Within distribution Evaluation

Given that DepthTrack [66], Lasher [36], and VisEvent [55] provide domain-specific testing sequences, our initial evaluation focuses on these within distribution sequences. For each dataset, we adhere to the metrics specified in the original papers and prior standards [69, 91] for evaluation.

**Comparison on DepthTrack [66]:** For evaluation, we use metrics such as precision (Pr) and recall (Re), as well as the F-score, which are the primary metrics. As shown in Tab. 1 When exclusively trained and fine-tuned on DepthTrack, Un-Track achieves a +2.1% absolute precision improvement over the current depth-specific SOTA ViPT [91]. Notably, even when trained with mixed data using a single parameter set, Un-Track outperforms the depth-specific ViPT. Furthermore, when the current ViPT is jointly trained on all datasets with a single set of parameters, its performance significantly deteriorates.

Additionally, we observe that trackers excelling in other tracking datasets [14, 22, 42] might struggle in RGB-D downstream settings. UniNext [65], a leading tracker trained on various large-scale tracking datasets and related image/ video tasks, exhibits poor performance. In contrast, our model achieves cross-modal unification within a single set of parameters, surpassing all depth-specific counterparts and performing closely to our specialized version. This underscores the efficacy of our shared embedding in achieving global alignment across diverse modalities.

	Thermal-Specific Parameters								Uni-model with a Single Set of Parameters					
	SGT [34]	FANet [96]	mfDiMP [81]	DAFNet [18]	MaCNet [77]	ProTrack [69]	ViPT [91]	Un-Track (ours)	Stark [62]	AiATrack [17]	OSTrack [72]	SeqTrack [8]	ViPT [91]	Un-Track (ours)
PR(↑)	0.327	0.441	0.447	0.448	0.482	0.538	<b>0.651</b>	<b>0.667</b>	0.418	0.463	0.530	0.582	0.608	<b>0.646</b>
SR(↑)	0.232	0.309	0.343	0.311	0.350	0.420	<b>0.525</b>	<b>0.536</b>	0.333	0.365	0.422	0.441	0.490	<b>0.513</b>

Table 2. Overall performance on the Lasher thermal testing set [36]. Our thermal-specific model sets a new SOTA record. Our uni-model variant competes strongly with the previous SOTA thermal-specific models and significantly surpasses its unimodel version.

	Event-Specific Parameters								Uni-model with a Single Set of Parameters					
	SiamCar [62]	Stark_E [62]	VITAL_E [46]	PrDiMP_E [13]	TransT_E [5]	ProTrack [69]	ViPT [91]	Un-Track (ours)	Stark [62]	AiATrack [17]	OSTrack [72]	SeqTrack [8]	ViPT [91]	Un-Track (ours)
Precision(↑)	0.599	0.612	0.649	0.644	0.650	0.632	<b>0.758</b>	<b>0.763</b>	0.616	0.626	0.691	0.665	0.740	<b>0.755</b>
Success(↑)	0.420	0.446	0.415	0.453	0.474	0.471	<b>0.592</b>	<b>0.597</b>	0.448	0.444	0.525	0.504	0.579	<b>0.589</b>

Table 3. Overall performance on VisEvent dataset [55]. Our event-specific model sets a new SOTA record. Our uni-variant, with the same parameter set as in Depth and Thermal, consistently achieves competitive performance across various modalities, leading to a significant margin over the uni-variant of the SOTA modality-specific model [91].

**Comparison on LasHer [36]:** Similarly, we conduct evaluations on the LasHer testing set for RGB-T tracker assessment as shown in Tab. 2. Precision (PR) and success rates (SR) are reported following conventional standards [36, 69, 91]. Initial comparisons under domain-specific settings reveal the challenges of achieving an architecture-unified model across RGB-D and RGB-T domains, with ProTrack [69] and ViPT [91] being the only works consistently leading in both settings. Our Un-Track, following domain-specific finetuning, outperforms the leading ViPT by a significant margin and sets a new SOTA record. In the cross-domain joint learning with a single set of parameters, our uni-model achieves a +3.8% absolute gain over ViPT. Remarkably, our model with a single set of parameters already achieves very competitive performance compared to the thermal-specific ViPT version.

**Comparison on VisEvent [55]:** We also evaluate tracker performance with RGB-Event input. Event data, being inherently sparse compared to depth and thermal information, presents challenges in extending existing RGB-D or RGB-T fusion designs for effective integration, hence necessitating specific fusion designs [3, 43, 98]. In contrast, we adopt a unified cross-modal prompting method based on shrinkage fusion. Our approach, with gradual token exchanges between RGB and event modalities, effectively preserves crucial modality-specific clues to enhance feature modeling. Performance-wise, under the event-specific setting, our Un-Track outperforms all counterparts, as shown in Tab. 3 and in Fig. 5. In the single set of parameters setting, our uni-model achieves a +1.1% absolute gain in precision over the current SOTA. This underscores the effectiveness of our approach in handling the unique challenges posed by event data integration, which can be mainly contributed to our shared binding that learns the global RGB+X alignment.




Figure 5. More precision/success comparisons on VisEvent dataset [55]. “Uni” stands for models with a single parameter set. “E” stands for the extension of RGB trackers with event fusion.

### 4.3. Generalization Across Datasets

In this section, we assess the versatility by evaluating performance on datasets that differ from the training ones, aligning with the goal of achieving a universal model checkpoint applicable to diverse scenarios.

**VOT-RGBD2022 [30]:** We first perform inference on the VOT-RGBD2022 dataset. Notably, our uni-model achieves superior performance compared to the depth-specific ViPT with a +0.5% absolute gain in accuracy.

**RGBT234 [35]:** We also test our model on the other thermal dataset RGBT234, which encompasses sequences with different distributions. Our uni-model surpasses the thermal-specific ViPT with a notable +0.7% absolute precision gain, as shown in Tab. 5.

Moreover, we present a detailed per-attribute comparison with the fine-tuned ViPT in Fig. 6. We are particularly interested in sequences with motion blur-fast motion-camera moving, as well as sequences with heavy occlusion-scale variants-background clutter. The former motion-related challenges can typically benefit from event clues, renowned for asynchronous computing, while the latter geometry-

	Depth-Specific Parameters								Uni-model with a Single Set of Parameters				
	DRefine [29]	Stark_D [29]	DMTracker [30]	DeT [66]	SBT_D [30]	SPT [94]	ProTrack [69]	ViPT [91]	Stark [62]	AiATrack [17]	OSTrack [72]	SeqTrack [8]	Un-Track (ours)
EAO(↑)	0.592	0.647	0.658	0.657	<b>0.708</b>	0.651	0.651	<b>0.721</b>	0.445	0.641	0.666	0.679	<b>0.718</b>
Accuracy(↑)	0.775	0.803	0.758	0.760	<b>0.809</b>	0.798	0.801	<b>0.815</b>	0.714	0.769	0.808	0.802	<b>0.820</b>
Robustness(↑)	0.760	0.798	0.851	0.845	<b>0.864</b>	0.851	0.802	<b>0.871</b>	0.598	0.832	0.814	0.846	<b>0.864</b>

Table 4. Overall performance on VOT-RGBD2022 [30]. Our uni-model, trained on a mix of all modalities, shows robust generalization and outperforms all depth-specific models and other uni-model counterparts.

	Thermal-Specific Parameters								Uni-model with a Single Set of Parameters						
	mfDiMP [81]	SGT [34]	DAFNet [18]	FANet [96]	MaCNet [77]	CMPP [52]	APFNet [59]	ProTrack [69]	ViPT [91]	Un-Track (ours)	Stark [62]	AiATrack [17]	OSTrack [72]	SeqTrack [8]	Un-Track (ours)
MPR(↑)	0.646	0.720	0.796	0.787	0.790	0.823	0.827	0.795	<b>0.835</b>	<b>0.837</b>	0.677	0.711	0.755	0.806	<b>0.842</b>
MSR(↑)	0.428	0.472	0.544	0.553	0.554	0.575	0.579	0.599	<b>0.617</b>	<b>0.618</b>	0.496	0.508	0.569	0.599	<b>0.625</b>

Table 5. Overall performance on RGBT234 dataset [35]. Our uni-model sets new SOTA records without specific thermal finetuning.




Figure 6. Per-attribute analysis on the thermal dataset RGBT234 [35]. Challenges related to motion and geometry are generally better addressed by event and depth cameras, respectively. Nevertheless, when inferring only with RGB-T data, our Un-Track surpasses both the SOTA thermal-specific method and the current leading uni-tracker. This success underscores our ability to learn emergent alignment across diverse modalities.

related challenges can typically benefit from depth cameras. However, these event/depth clues are not directly available in the RGB-T setting. Nevertheless, our Un-Track, trained on all modalities, outperforms both thermal fine-tuned ViPT [91] and the current leading uni-tracker [8] with large margins. This underscores our capability in learning event and depth priors through shared binding, without the need for the presence of these modalities during inference.

**RGB-only:** In practical scenarios, challenges arise when there are no modal clues available, a typical case is when the auxiliary sensor fails to work properly. We address this demanding case in our study by substituting the modal input with dummy values. Under a such challenging setting, as shown in Tab. 6, our uni-method consistently outperforms both our RGB baseline fine-tuned counterparts significantly. Notably, such an improvement is achieved with a very limited increase in learning parameters, *i.e.*, +6.65M with +2.14 GFLOPs.

	RGB Baseline [72]	ViPT [91]	Un-Track (ours)
GFLOPs	21.50	21.80	23.64
Params (M)	92.08	92.96	98.73
F-score(↑)	0.529	0.542	0.558
Re(↑)	0.522	0.538	0.557
Pr(↑)	0.536	0.546	0.560

Table 6. Overall performance on DepthTrack test set [66] with dummy depth input.

	w/o Shared Embed	[91] as Modal Prompt	w/o LoRA Finetune
F-score(↑)	0.599	0.579	0.594
Re(↑)	0.602	0.575	0.598
Pr(↑)	0.597	0.584	0.596

Table 7. Key component analysis.

## 5. Ablation Studies

In this section, we perform all experiments on the DepthTrack testing set [66] under a single parameter set setting.

**Key Component Analysis:** We begin by studying the effectiveness of key components, including the shared embedding, modal prompting, and LoRA finetuning, as summarized in Tab. 7. We initially remove the shared embedding by directly feeding mixed modalities into the learning diagram. It can be seen that this equal treatment of all modalities harms network performance due to the heterogeneous representation across modal domains. Secondly, we replace our prompting block with a recent counterpart that computes fovea attention from additional input [91]. Our designed gradual shrinkage fusion, allowing token-wise interaction, proves to be more effective. We also report performance when we remove the inner fine-tuning with Lora. It can be seen that the performance deteriorates.

**Low-Rank:** Low-rank approximation plays a vital role in our model, influencing our shared embedding, modal fusion, and LoRA-finetuning. Hence, the choice of rank is crucial for both the efficiency and effectiveness of our ap-

Table 8. **Low Rank Approximation** Our plain version is highlighted in gray.

	(a) $Rank_k$ (Sec. 3.2)			(b) $Rank_l$ (Sec. 3.3)			(c) LoRA (Sec. 3.4)			(d) Percentile (Sec. 3.3)		
	2	4	8	4	8	16	2	4	8	1/8	1/4	1/3
F-score( $\uparrow$ )	0.607	0.610	0.602	0.596	0.610	0.606	0.601	0.610	0.600	0.604	0.610	0.595
Re( $\uparrow$ )	0.606	0.608	0.601	0.593	0.608	0.609	0.599	0.608	0.598	0.606	0.608	0.593
Pr( $\uparrow$ )	0.608	0.611	0.604	0.599	0.611	0.604	0.602	0.611	0.602	0.602	0.611	0.596

proach. In Tab. 8, we systematically explore the impact of low-rank choices within each component.

**Shared Embedding:** Our objective is to identify an optimal low-rank latent space for merging different modalities effectively. Tab. 8(a) presents our ablation study, where we explore ranks of 2, 4, and 8. Lower ranks result in poorer performance due to sparse representations. Conversely, higher ranks capture too many modality-specific details, complicating the search for a shared embedding.

**Modal Prompting:** As shown in Tab. 8(b), similar trends are observed when investigating the low-rank choices for modal prompting with rank values of 4, 8, and 16, as low ranks struggle to capture essential information, while higher ranks introduce an overload of modality-specific details.

**LoRA-finetuning:** We also vary the ranks between 2, 4, and 8 for the LoRA finetuning technique, as shown in Tab. 8(c). Lower ranks exhibited consistently poor performance, while higher ranks in this case tend to result in poorer performance, likely due to overfitting.

**Remarks:** These experiments emphasize the importance of selecting the best low-rank representations. Nevertheless, our model shows great resilience to the choice of LoRA, showcasing consistent performance across different low-rank configurations. Notably, all our low-rank variants outperform the current SOTA ViPT under the uni-setting, validating our robustness and effectiveness.

**Modal Prompting:** During the prompting fusion, a learnable score function is used to categorize tokens into three groups based on their confidence scores. Here, we explore different percentiles for the number of positive, which is the same as the number of negative tokens, leaving the rest as uncertain tokens. As shown in Tab. 8(d), the choice of percentile can influence overall performance. Lower percentiles result in poorer performance, as recovering uncertain tokens from very few neighboring tokens is challenging. Higher percentiles also lead to performance degradation since the focus is on token exchange rather than token fusion. The choice of 1/4 leads to the best balance between exchanger and fuser, leading to the best performance.

**Shared Embedding:** Here, we perform ablation studies on our shared embedding, a foundational component of our uni-model. The quantitative results are presented in Tab. 9. We begin by exploring a scenario where our shared embedding is replaced with a variant lacking explicit edge guidance (w/o Explicit Edge). In this configuration, the network

	w/o Explicit Edge	w/o Implicit Learning	w/o In-domain Approx.
F-score( $\uparrow$ )	0.600	0.604	0.581
Re( $\uparrow$ )	0.602	0.609	0.583
Pr( $\uparrow$ )	0.598	0.599	0.579

Table 9. Ablation on shared embedding

learns the shared embedding solely without any edge prior. The results highlight a substantial performance drop, underscoring the pivotal role of explicit edge guidance — a natural and static embedding that binds all modalities together — in facilitating this implicit learning process.

We also conduct experiments using only the explicit edge as the shared embedding, excluding any additional learning modules (w/o Implicit Learning). This approach, too, yields suboptimal performance, primarily due to the neglect of modality-specific clues within each domain.

Finally, we directly compute the low-rank approximation from mixed modalities, bypassing the initial in-domain approximation and subsequent fusion steps (w/o In-domain Approx.). We observe that the network struggles in learning a shared embedding in this direct low-rank approximation setup, primarily due to the intricately mixed representation caused by the domain gap between different modalities.

## 6. Conclusion and Future Work

We present a successful case of a single-model and any-modality tracker for video object tracking. The proposed method achieves a shared embedding that binds all modalities together, overcoming their heterogeneous representations. This unification is facilitated by lightweight modal prompting and inner finetuning, inheriting benefits from large-scale pre-trained trackers without introducing a substantial computational burden. Exhaustive experiments showcase our improved tracking performance and robust generalization, with any modality input.

**Acknowledgements.** The authors thank the reviewers and ACs for their tremendous efforts and helpful comments. The event icon is credited to Zuowen Wang. This research is financed in part by the Alexander von Humboldt Foundation, in part by NSFC (62376156, 62322113), and in part by the Ministry of Education and Science of Bulgaria (support for INSAIT, part of the Bulgarian National Roadmap for Research Infrastructure).



## References

- [1] Simon Arridge, Pascal Fensel, and Andreas Hauptmann. Joint reconstruction and low-rank decomposition for dynamic inverse problems. *Inverse Problems and Imaging*, 16(3):483–523, 2022. [2](#)
- [2] Boyu Chen, Peixia Li, Lei Bai, Lei Qiao, Qihong Shen, Bo Li, Weihao Gan, Wei Wu, and Wanli Ouyang. Backbone is all your need: A simplified architecture for visual object tracking. In *ECCV*, 2022. [2](#)
- [3] Qinyu Chen, Zuowen Wang, Shih-Chii Liu, and Chang Gao. 3et: Efficient event-based eye tracking using a change-based convlstm network. *arXiv preprint arXiv:2308.11771*, 2023. [6](#)
- [4] Wanli Chen, Xinge Zhu, Ruoqi Sun, Junjun He, Ruiyu Li, Xiaoyong Shen, and Bei Yu. Tensor low-rank reconstruction for semantic segmentation. In *ECCV*, 2020. [2](#)
- [5] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *CVPR*, 2021. [2](#), [6](#)
- [6] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *CVPR*, 2021. [2](#)
- [7] Xin Chen, Bin Yan, Jiawen Zhu, Huchuan Lu, Xiang Ruan, and Dong Wang. High-performance transformer tracking. *TPAMI*, 2022. [2](#)
- [8] Xin Chen, Houwen Peng, Dong Wang, Huchuan Lu, and Han Hu. Seqtrack: Sequence to sequence learning for visual object tracking. In *CVPR*, 2023. [5](#), [6](#), [7](#)
- [9] Zedu Chen, Bineng Zhong, Guorong Li, Shengping Zhang, and Rongrong Ji. Siamese box adaptive network for visual tracking. In *CVPR*, 2020. [1](#)
- [10] Anthony Cioppa, Silvio Giancola, Adrien Deliege, Le Kang, Xin Zhou, Zhiyu Cheng, Bernard Ghanem, and Marc Van Droogenbroeck. Soccernet-tracking: Multiple object tracking dataset and benchmark in soccer videos. In *CVPR*, 2022. [1](#)
- [11] Yutao Cui, Cheng Jiang, Limin Wang, and Gangshan Wu. Mixformer: End-to-end tracking with iterative mixed attention. In *CVPR*, 2022. [2](#)
- [12] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. ATOM: Accurate tracking by overlap maximization. In *CVPR*, 2019. [1](#)
- [13] Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *CVPR*, 2020. [6](#)
- [14] 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. [2](#), [5](#)
- [15] Yingkai Fu, Meng Li, Wenxi Liu, Yuanchen Wang, Jiqing Zhang, Baocai Yin, Xiaopeng Wei, and Xin Yang. Distractor-aware event-based tracking. *TIP*, 2023. [1](#)
- [16] Shang Gao, Jinyu Yang, Zhe Li, Feng Zheng, Aleš Leonardis, and Jingkuan Song. Learning dual-fused modality-aware representations for rgb-d tracking. In *ECCV*, 2022. [2](#)
- [17] Shenyuan Gao, Chunluan Zhou, Chao Ma, Xinggang Wang, and Junsong Yuan. Aiatrack: Attention in attention for transformer visual tracking. In *ECCV*, 2022. [5](#), [6](#), [7](#)
- [18] Yuan Gao, Chenglong Li, Yabin Zhu, Jin Tang, Tao He, and Futian Wang. Deep adaptive fusion network for high performance RGB-D tracking. In *ICCVW*, 2019. [6](#), [7](#)
- [19] Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *CVPR*, 2023. [2](#)
- [20] Botao He, Haojia Li, Siyuan Wu, Dong Wang, Zhiwei Zhang, Qianli Dong, Chao Xu, and Fei Gao. Fast-dynamic-vision: Detection and tracking dynamic objects with event and depth sensing. In *IROS*, 2021. [1](#)
- [21] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022. [3](#), [4](#), [5](#)
- [22] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *TPAMI*, 2019. [2](#), [5](#)
- [23] Sajid Javed, Martin Danelljan, Fahad Shahbaz Khan, Muhammad Haris Khan, Michael Felsberg, and Jiri Matas. Visual object tracking with discriminative filters and siamese networks: a survey and outlook. *TPAMI*, 45(5):6552–6574, 2022. [1](#)
- [24] I-Hong Jhuo, Dong Liu, DT Lee, and Shih-Fu Chang. Robust visual domain adaptation with low-rank reconstruction. In *CVPR*, 2012. [2](#)
- [25] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, 2022. [2](#), [4](#)
- [26] Ben Kang, Xin Chen, Dong Wang, Houwen Peng, and Huchuan Lu. Exploring lightweight hierarchical vision transformers for efficient visual tracking. In *ICCV*, 2023. [2](#)
- [27] Muhammad Uzair Khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *CVPR*, 2023. [2](#)
- [28] 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. [5](#)
- [29] Matej Kristan, Jiří Matas, Aleš Leonardis, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Hyung Jin Chang, Martin Danelljan, Luka Čehovin, Alan Lukežič, et al. The ninth visual object tracking vot2021 challenge results. In *ICCVW*, 2021. [7](#)
- [30] 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](#), [7](#)
- [31] Yi-Lun Lee, Yi-Hsuan Tsai, Wei-Chen Chiu, and Chen-Yu Lee. Multimodal prompting with missing modalities for visual recognition. In *CVPR*, 2023. [2](#), [3](#)
- [32] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *CVPR*, 2018. [2](#)

- [33] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. SiamRPN++: Evolution of siamese visual tracking with very deep networks. In *CVPR*, 2019. 2
- [34] Chenglong Li, Nan Zhao, Yijuan Lu, Chengli Zhu, and Jin Tang. Weighted sparse representation regularized graph learning for RGB-T object tracking. In *ACM MM*, 2017. 6, 7
- [35] Chenglong Li, Xinyan Liang, Yijuan Lu, Nan Zhao, and Jin Tang. RGB-T object tracking: Benchmark and baseline. *PR*, 96:106977, 2019. 2, 6, 7
- [36] 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. 2, 5, 6
- [37] Qiao Liu, Xin Li, Zhenyu He, Nana Fan, Di Yuan, and Hongpeng Wang. Learning deep multi-level similarity for thermal infrared object tracking. *TMM*, 23:2114–2126, 2020. 2
- [38] Cheng Long Li, Andong Lu, Ai Hua Zheng, Zhengzheng Tu, and Jin Tang. Multi-adaptor rgbt tracking. In *ICCVW*, 2019. 2
- [39] Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Motlaghi, and Aniruddha Kembhavi. Unified-io: A unified model for vision, language, and multi-modal tasks. In *ICLR*, 2023. 2
- [40] Alan Lukezic, Ugur Kart, Jani Kapyla, Ahmed Durmush, Joni-Kristian Kamarainen, Jiri Matas, and Matej Kristan. Cdtb: A color and depth visual object tracking dataset and benchmark. In *ICCV*, 2019. 1
- [41] Mengmeng Ma, Jian Ren, Long Zhao, Sergey Tulyakov, Cathy Wu, and Xi Peng. Smil: Multimodal learning with severely missing modality. In *AAAI*, 2021. 2
- [42] 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. 2, 5
- [43] Yansong Peng, Yueyi Zhang, Zhiwei Xiong, Xiaoyan Sun, and Feng Wu. Get: Group event transformer for event-based vision. In *ICCV*, 2023. 6
- [44] Yansheng Qiu, Ziyuan Zhao, Hongdou Yao, Delin Chen, and Zheng Wang. Modal-aware visual prompting for incomplete multi-modal brain tumor segmentation. In *ACM MM*, 2023. 2
- [45] Esteban Real, Jonathon Shlens, Stefano Mazzocchi, Xin Pan, and Vincent Vanhoucke. Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video. In *CVPR*, 2017. 2
- [46] Yibing Song, Chao Ma, Xiaohe 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. 6
- [47] Daniel Stadler and Jurgen Beyerer. Improving multiple pedestrian tracking by track management and occlusion handling. In *CVPR*, 2021. 1
- [48] Chuanming Tang, Xiao Wang, Ju Huang, Bo Jiang, Lin Zhu, Jianlin Zhang, Yaowei Wang, and Yonghong Tian. Revisiting color-event based tracking: A unified network, dataset, and metric. *arXiv preprint arXiv:2211.11010*, 2022. 1
- [49] Chuanming Tang, Xiao Wang, Ju Huang, Bo Jiang, Lin Zhu, Jianlin Zhang, Yaowei Wang, and Yonghong Tian. Revisiting color-event based tracking: A unified network, dataset, and metric. *arXiv preprint arXiv:2211.11010*, 2022. 2
- [50] Paul Voigtlaender, Jonathon Luiten, Philip H. S. Torr, and Bastian Leibe. Siam R-CNN: Visual tracking by re-detection. In *CVPR*, 2020. 2
- [51] Zhexiong Wan, Yuxin Mao, Jing Zhang, and Yuchao Dai. Rpeflow: Multimodal fusion of rgb-pointcloud-event for joint optical flow and scene flow estimation. In *ICCV*, 2023. 2
- [52] Chaoqun Wang, Chunyan Xu, Zhen Cui, Ling Zhou, Tong Zhang, Xiaoya Zhang, and Jian Yang. Cross-modal pattern-propagation for RGB-T tracking. In *CVPR*, 2020. 7
- [53] Hu Wang, Yuanhong Chen, Congbo Ma, Jodie Avery, Louise Hull, and Gustavo Carneiro. Multi-modal learning with missing modality via shared-specific feature modelling. In *CVPR*, 2023. 3
- [54] Ning Wang, Wengang Zhou, Jie Wang, and Houqiang Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *CVPR*, 2021. 2
- [55] 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. *arXiv preprint arXiv:2108.05015*, 2021. 2, 5, 6
- [56] Xiao Wang, Xiujun Shu, Shilliang Zhang, Bo Jiang, Yaowei Wang, Yonghong Tian, and Feng Wu. Mfgnet: Dynamic modality-aware filter generation for rgb-t tracking. *TMM*, 2022. 2
- [57] Zuowen Wang, Yuhuang Hu, and Shih-Chii Liu. Exploiting spatial sparsity for event cameras with visual transformers. In *ICIP*, 2022. 2
- [58] David Wisth, Marco Camurri, Sandipan Das, and Maurice Fallon. Unified multi-modal landmark tracking for tightly coupled lidar-visual-inertial odometry. *RA-L*, 6(2):1004–1011, 2021. 2
- [59] Yun Xiao, Mengmeng Yang, Chenglong Li, Lei Liu, and Jin Tang. Attribute-based progressive fusion network for RGBT tracking. In *AAAI*, 2022. 7
- [60] Yun Xiao, Mengmeng Yang, Chenglong Li, Lei Liu, and Jin Tang. Attribute-based progressive fusion network for rgbt tracking. In *AAAI*, 2022. 1, 2
- [61] Yinda Xu, Zeyu Wang, Zuoxin Li, Ye Yuan, and Gang Yu. SiamFC++: Towards robust and accurate visual tracking with target estimation guidelines. In *AAAI*, 2020. 2
- [62] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *ICCV*, 2021. 5, 6, 7
- [63] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *ICCV*, 2021. 2
- [64] Bin Yan, Houwen Peng, Kan Wu, Dong Wang, Jianlong Fu, and Huchuan Lu. Lighttrack: Finding lightweight neural networks for object tracking via one-shot architecture search. In *CVPR*, 2021. 2
- [65] Bin Yan, Yi Jiang, Jiannan Wu, Dong Wang, Ping Luo, Zehuan Yuan, and Huchuan Lu. Universal instance perception as object discovery and retrieval. In *CVPR*, 2023. 5

- [66] Song Yan, Jinyu Yang, Jani Käpylä, Feng Zheng, Aleš Leonardis, and Joni-Kristian Kämäräinen. Depthtrack: Unveiling the power of RGBD tracking. In *ICCV*, 2021. 2, 5, 7
- [67] Song Yan, Jinyu Yang, Ales Leonardis, and Joni-Kristian Kamarainen. Depth-only object tracking. In *BMVC*, 2021. 2
- [68] Jinyu Yang, Zhe Li, Song Yan, Feng Zheng, Aleš Leonardis, Joni-Kristian Kämäräinen, and Ling Shao. Rgbd object tracking: An in-depth review. *arXiv preprint arXiv:2203.14134*, 2022. 1
- [69] Jinyu Yang, Zhe Li, Feng Zheng, Ales Leonardis, and Jingkuan Song. Prompting for multi-modal tracking. In *ACMMM*, 2022. 1, 2, 5, 6, 7
- [70] Jinyu Yang, Shang Gao, Zhe Li, Feng Zheng, and Aleš Leonardis. Resource-efficient rgbd aerial tracking. In *CVPR*, 2023. 2
- [71] Rui Yao, Guosheng Lin, Shixiong Xia, Jiaqi Zhao, and Yong Zhou. Video object segmentation and tracking: A survey. *ACM TIST*, 11(4):1–47, 2020. 2
- [72] 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. 3, 5, 6, 7
- [73] Jie Yin, Ang Li, Tao Li, Wenxian Yu, and Danping Zou. M2dgr: A multi-sensor and multi-scenario slam dataset for ground robots. *RA-L*, 7(2):2266–2273, 2021. 2
- [74] Yuechen Yu, Yilei Xiong, Weilin Huang, and Matthew R Scott. Deformable siamese attention networks for visual object tracking. In *CVPR*, 2020. 2
- [75] Jiandian Zeng, Tianyi Liu, and Jiantao Zhou. Tag-assisted multimodal sentiment analysis under uncertain missing modalities. In *ACM SIGIR*, 2022. 2
- [76] 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*, 2023. 2
- [77] Hui Zhang, Lei Zhang, Li Zhuo, and Jing Zhang. Object tracking in RGB-T videos using modal-aware attention network and competitive learning. *Sensors*, 20(2):393, 2020. 6, 7
- [78] Jiqing Zhang, Xin Yang, Yingkai Fu, Xiaopeng Wei, Baocai Yin, and Bo Dong. Object tracking by jointly exploiting frame and event domain. In *ICCV*, 2021. 2
- [79] Jiqing Zhang, Bo Dong, Haiwei Zhang, Jianchuan Ding, Felix Heide, Baocai Yin, and Xin Yang. Spiking transformers for event-based single object tracking. In *CVPR*, 2022. 1
- [80] Jiaming Zhang, Ruiping Liu, Hao Shi, Kailun Yang, Simon Reiß, Kunyu Peng, Haodong Fu, Kaiwei Wang, and Rainer Stiefelhagen. Delivering arbitrary-modal semantic segmentation. In *CVPR*, 2023. 2
- [81] Lichao Zhang, Martin Danelljan, Abel Gonzalez-Garcia, Joost van de Weijer, and Fahad Shahbaz Khan. Multi-modal fusion for end-to-end RGB-T tracking. In *ICCVW*, 2019. 6, 7
- [82] Pengyu Zhang, Dong Wang, and Huchuan Lu. Multi-modal visual tracking: Review and experimental comparison. *arXiv preprint arXiv:2012.04176*, 2020. 1
- [83] Pengyu Zhang, Jie Zhao, Dong Wang, Huchuan Lu, and Xiang Ruan. Visible-thermal uav tracking: A large-scale benchmark and new baseline. In *CVPR*, 2022. 1
- [84] Wenwei Zhang, Hui Zhou, Shuyang Sun, Zhe Wang, Jianping Shi, and Chen Change Loy. Robust multi-modality multi-object tracking. In *ICCV*, 2019. 1
- [85] Zhipeng Zhang and Houwen Peng. Deeper and wider siamese networks for real-time visual tracking. In *CVPR*, 2019. 2
- [86] Haojie Zhao, Dong Wang, and Huchuan Lu. Representation learning for visual object tracking by masked appearance transfer. In *CVPR*, 2023. 1
- [87] Jinjian Zhao, Xiaohan Zhang, and Pengyu Zhang. A unified approach for tracking uavs in infrared. In *ICCV*, 2021. 1
- [88] Shaochuan Zhao, Tianyang Xu, Xiao-Jun Wu, and Xue-Feng Zhu. Adaptive feature fusion for visual object tracking. *PR*, 111:107679, 2021. 2
- [89] Aihua Zheng, Zi Wang, Zihan Chen, Chenglong Li, and Jin Tang. Robust multi-modality person re-identification. In *AAAI*, 2021. 1
- [90] Jiawen Zhu, Zhenyu Chen, Zeqi Hao, Shijie Chang, Lu Zhang, Dong Wang, Huchuan Lu, Bin Luo, Jun-Yan He, Jin-Peng Lan, et al. Tracking anything in high quality. *arXiv preprint arXiv:2307.13974*, 2023. 1
- [91] Jiawen Zhu, Simiao Lai, Xin Chen, Dong Wang, and Huchuan Lu. Visual prompt multi-modal tracking. In *CVPR*, 2023. 1, 2, 5, 6, 7
- [92] Jiawen Zhu, Huayi Tang, Zhi-Qi Cheng, Jun-Yan He, Bin Luo, Shihao Qiu, Shengming Li, and Huchuan Lu. Dcpt: Darkness clue-prompted tracking in nighttime uavs. *arXiv preprint arXiv:2309.10491*, 2023. 1
- [93] Xue-Feng Zhu, Xiao-Jun Wu, Tianyang Xu, Zhen-Hua Feng, and Josef Kittler. Robust visual object tracking via adaptive attribute-aware discriminative correlation filters. *TMM*, 24:301–312, 2021. 2
- [94] 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. *AAAI*, 2023. 2, 5, 7
- [95] 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
- [96] Yabin Zhu, Chenglong Li, Jin Tang, and Bin Luo. Quality-aware feature aggregation network for robust RGBT tracking. *IEEE TIV*, 6(1):121–130, 2020. 6, 7
- [97] Zhiyu Zhu, Junhui Hou, and Xianqiang Lyu. Learning graph-embedded key-event back-tracing for object tracking in event clouds. *NeurIPS*, 2022. 2
- [98] Zhiyu Zhu, Junhui Hou, and Dapeng Oliver Wu. Cross-modal orthogonal high-rank augmentation for rgb-event transformer-trackers. In *ICCV*, 2023. 1, 6