

Visual Prompt Multi-Modal Tracking

Jiawen Zhu^{1†} Simiao Lai^{1†} Xin Chen¹ Dong Wang^{1✉} Huchuan Lu^{1,2}
¹Dalian University of Technology, China ²Peng Cheng Laboratory, China

{jiawen, laisimiao, chenxin3131}@mail.dlut.edu.cn, {wdice, lhchuan}@dlut.edu.cn

Abstract

Visible-modal object tracking gives rise to a series of downstream multi-modal tracking tributaries. To inherit the powerful representations of the foundation model, a natural *modus operandi* for multi-modal tracking is full fine-tuning on the RGB-based parameters. Albeit effective, this manner is not optimal due to the scarcity of downstream data and poor transferability, etc. In this paper, inspired by the recent success of the prompt learning in language models, we develop **Visual Prompt multi-modal Tracking (ViPT)**, which learns the modal-relevant prompts to adapt the frozen pre-trained foundation model to various downstream multi-modal tracking tasks. ViPT finds a better way to stimulate the knowledge of the RGB-based model that is pre-trained at scale, meanwhile only introducing a few trainable parameters (less than 1% of model parameters). ViPT outperforms the full fine-tuning paradigm on multiple downstream tracking tasks including RGB+Depth, RGB+Thermal, and RGB+Event tracking. Extensive experiments show the potential of visual prompt learning for multi-modal tracking, and ViPT can achieve state-of-the-art performance while satisfying parameter efficiency. Code and models are available at <https://github.com/jiawen-zhu/ViPT>.

1. Introduction

RGB-based tracking, a foundation task of visual object tracking, gains from large-scale benchmarks [9, 13, 27, 36, 39, 45] provided by the community, and many excellent works [2, 3, 5, 7, 22] have spurred out over the past decades. Despite the promising results, object tracking based on pure RGB sequences is still prone to failure in some complex and corner scenarios, *e.g.*, extreme illumination, background clutter, and motion blur. Therefore, multi-modal tracking is drawing increasing attention due to the ability to achieve more robust tracking by utilizing inter-modal complementarity, among which RGB+Depth (RGB-D) [47, 60],

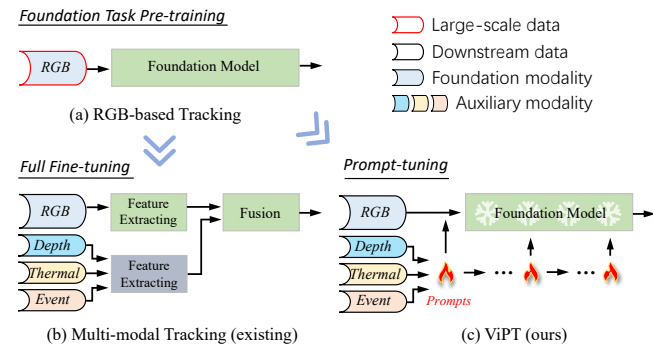


Figure 1. **Existing multi-modal tracking paradigm vs. ViPT.** (a) Foundation model training on large-scale RGB sequences. (a)→(b) Existing multi-modal methods extend the off-the-shelf foundation model and conduct full fine-tuning on downstream tracking tasks. (a)→(c) We investigate the design of a multi-modal tracking method in prompt-learning paradigm. Compared to (b), the proposed ViPT has a more concise network structure, benefits from parameter-friendly prompt-tuning, and narrows the gap between the foundation and downstream models.

RGB+Thermal (RGB-T) [44, 56], and RGB+Event (RGB-E) [51, 52] are represented.

However, as the downstream task of RGB-based tracking, the main issue encountered by multi-modal tracking is the lack of large-scale datasets. For example, the widely used RGB-based tracking datasets, GOT-10k [13], TrackingNet [36], and LaSOT [9], contain 9.3K, 30.1K, and 1.1K sequences, corresponding to 1.4M, 14M, and 2.8M frames for training. Whereas the largest training datasets in multi-modal tracking, DepthTrack [47], LasHeR [25], VisEvent [43], contain 150, 979, 500 training sequences, corresponding to 0.22M, 0.51M, 0.21M annotated frame pairs, which is at least an order of magnitude less than the former. Accounting for the above limitation, multi-modal tracking methods [43, 47, 61] usually utilize pre-trained RGB-based trackers and perform fine-tuning on their task-oriented training sets (as shown in Figure 1 (a)→(b)). DeT [47] adds a depth feature extraction branch to the original ATOM [7] or DiMP [3] tracker and fine-tunes on RGB-D training data. Zhang *et al.* [57] extend SiamRPN++ [21] with dual-modal inputs for RGB-T tracking. They first con-

[†]Equal contribution.

[✉]Corresponding author: Dr. Dong Wang.

struct a unimodal tracking network trained on RGB data, then tune the whole extended multi-modal network with RGB-T image pairs. Similarly, Wang *et al.* [43] develop dual-modal trackers by extending single-modal ones with various fusion strategies for visible and event flows and perform extra model training on the RGB-E sequences. Although effective, the task-oriented full-tuning approach has some drawbacks. (i) *Full fine-tuning the model is time expensive and inefficient, and the burden of parameters storing is large, which is unfriendly to numerous applications and cumbersome to transfer deploying.* (ii) *Full fine-tuning is unable to obtain generalized representation due to the limited annotated samples, the inability to utilize the pre-trained knowledge of the foundation model trained on large-scale datasets.* Thus, a natural question is thrown up: Is there a more effective manner to adapt the RGB-based foundation model to downstream multi-modal tracking?

More recently, in Natural Language Processing (NLP) field, researchers have injected textual prompts into downstream language models to effectively exploit the representational potential of foundation models, this method, is known as prompt-tuning. After that, a few researchers have tried to freeze the entire upstream model and add only some learnable parameters to the input side to learn valid visual prompts. Existing researches [1, 14, 38, 59] show its great potential and visual prompt learning is expected to be an alternative to full fine-tuning. Intuitively, there is a large inheritance between multi-modal and single RGB-modal tracking, which should share most of prior knowledge on feature extraction or attention patterns. In this spirit, we present *ViPT*, a unified visual prompt-tuning paradigm for downstream multi-modal tracking. Instead of fully fine-tuning an RGB-based tracker combined with an auxiliary-modal branch, *ViPT* freezes the whole foundation model and only learns a few modal-specific visual prompts, which inherits the RGB-based model parameters trained at scale to the maximum extent (see Figure 1 (c)). Different from the prompt learning of other single-modal vision tasks, *ViPT* introduces additional auxiliary-modal inputs into the prompt-tuning process, adapting the foundation model to downstream tasks while simultaneously learning the association between different modalities. Specifically, *ViPT* inserts several simple and lightweight modality-complementary prompter (MCP) blocks into the frozen foundation model to effectively learn the inter-modal complementarities. Notably, *ViPT* is a general framework for various downstream multi-modal tracking tasks, including RGB-D, RGB-T, and RGB-E tracking. We summarize the contribution of our work as follows:

- A visual prompt tracking framework is proposed to achieve task-oriented multi-modal tracking. Facilitated by learned prompts, the off-the-shelf foundation model can be effectively adapted from RGB domain

to downstream multi-modal tracking tasks. Besides, *ViPT* is a general method that can be applied to various tasks, *i.e.*, RGB-D, RGB-T, and RGB-E tracking.

- A modality-complementary prompter is designed to generate valid visual prompts for the task-oriented multi-modal tracking. The auxiliary-modal inputs are streamlined to a small number of prompts instead of designing an extra network branch.
- Extensive experiments show that our method achieves SOTA performance on multiple downstream multi-modal tracking tasks while maintaining parameter-efficient (<1% trainable parameters).

2. Related Work

2.1. Multi-Modal tracking

Assigned an arbitrary object in the initial frame, general visual object tracking (single-modal inputs) aims to predict the position and scale of the object in all subsequent frames. The continued emergence of large-scale datasets [9, 13, 36], and the surging development of deep neural networks [12, 41] have further advanced this field. Numerous trackers [2, 3, 5, 22] have been deployed to refresh the accuracy and robustness of tracking. Even though the existing performance is already impressive, single-modal object tracking in some scenarios (*e.g.*, background clutter, illumination variation, motion blur, *etc.*) is still insufficient to meet our demands. As the cost of binocular imaging sensors becomes increasingly low, multi-modal tracking has attracted more attention for it can provide complementary information among different modalities, assisting trackers to handle challenging situations that cannot be addressed solely through RGB inputs. For example, [35] assists with additional depth information for hand tracking in occlusion scenes. Similarly, JMMAC [55] fuses two modalities, RGB and thermal infrared, to obtain reliable appearance and motion cues. To obtain more robust tracking results, Liu *et al.* [28] combines RGB and event flows to predict candidate ROIs in complicated scenarios. In addition, numerous works focus on how to perform effective feature fusion on multi-modal information [6, 33, 52, 58].

However, an issue has been overlooked. A potential risk of overfitting is rising due to the contradiction between the paucity of large-scale multi-modal training sets and the huge appetites of the data-driven models. Therefore, existing methods [31, 43, 47, 53, 54, 60, 61] usually design an additional network branch for the inputs of the auxiliary modality, and the final model first loads the pre-trained parameters (from RGB-based model or imagenet [39]) then conducts fine-tuning on their downstream task-oriented training sets. In addition, Zhang *et al.* [53] transfer the RGB videos in GOT10k [13] to synthetic TIR videos, and Yan *et al.* [47] generate synthetic RGB-D data from LaSOT [9]

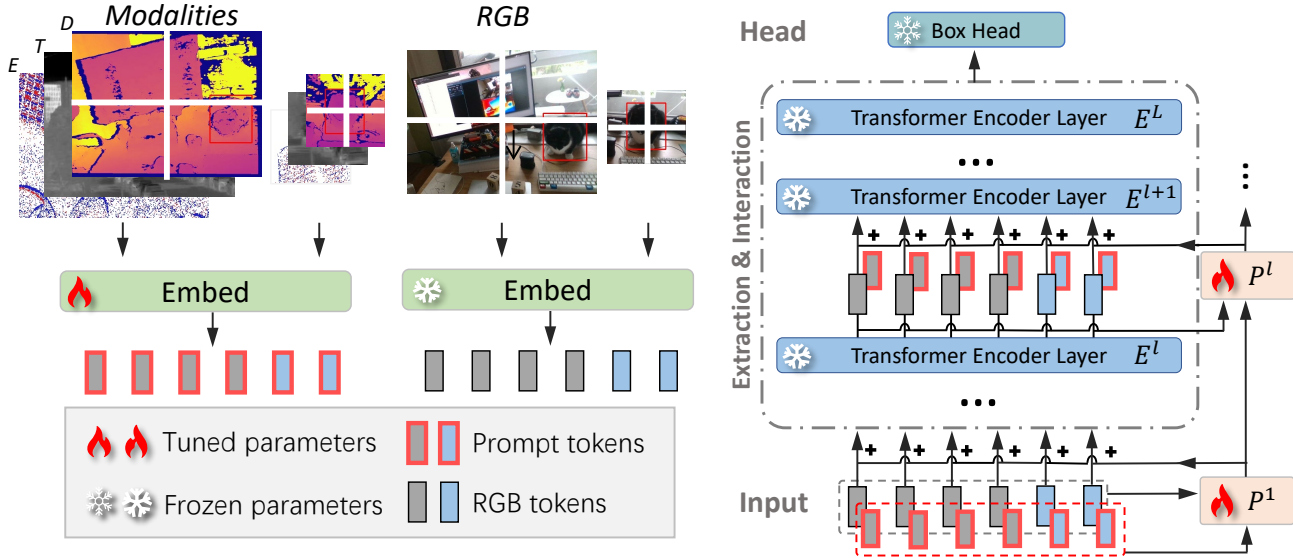


Figure 2. **Overview architecture of our ViPT.** The RGB- and auxiliary- modal inputs are first fed to the patch embed to generate the corresponding RGB and prompt tokens. L -layer stacked vision transformer (ViT) backbone is used for feature extraction and interaction. Modality-complementary prompts $\{P^l\}_{l=1}^L$ are inserted into the foundation model to learn the effective visual prompts.

and COCO [27]. They attempted to exploit some model-transform algorithms (e.g., [34] for monocular depth estimation) to generate auxiliary-modal pseudo-truths from RGB sequences to expand the downstream data, whereas this approach just partly alleviates the data anxiety for multi-modal tracking. In this work, taking a different path, we design a parameter-efficient prompt-tuning architecture to explore the adaptation from pre-trained RGB-based model to downstream multi-modal tracking tasks.

2.2. Visual Prompt Learning

For a long period, fine-tuning has been adopted to leverage large pre-trained models for performing downstream tasks. Researchers usually update all the model parameters when training on downstream data. This manner is parameter-inefficient and requires many repetitive task-oriented copies and storage of the entire pre-trained model. Recently, as a new paradigm, prompt learning has emerged and drastically boosted the performance of many downstream natural language processing (NLP) tasks [20, 30]. Besides, prompt learning also shows its effectiveness in many computer vision tasks. VPT [14] prepends a set of learnable parameters to transformer encoders and remarkably beats full fine-tuning on 20 downstream recognition tasks. AdaptFormer [4] introduces lightweight modules to ViT and outperforms full fine-tuned models on action recognition benchmarks. Convpass [15] is proposed to employ convolutional bypasses for prompt learning on pre-trained ViT. ProTrack [48] first introduces the prompting concept into the tracking field. However, it merely simply

fuses source images from two modalities by weighted addition without the tuning process. The prompts in it are not learnable and thus can gain limited improvements. In this work, we exploit the core idea of visual prompt-tuning and design an innovative prompt-learning framework for multi-modal tracking to replace the full fine-tuning paradigm.

3. Methodology

In this work, we propose ViPT for effectively and efficiently adapting the pre-trained RGB-based foundation tracking model to downstream multi-modal tasks. Instead of fully fine-tuning upon a pre-trained foundation model, ViPT only tunes a small number of prompt-learning parameters and achieves promising transfer learning abilities and preferable modal complementarities. The overall architecture of our ViPT is depicted in Figure 2.

3.1. Preliminary and Notation

Problem Setting. Given a video with an initial target box B_0 , the goal of the RGB-based tracking is to learn a tracker $F_{RGB} : \{X_{RGB}, B_0\} \rightarrow B$ to predict the box B of the target in all subsequent search frames X_{RGB} . For multi-modal tracking, it introduces an extra spatial-temporal synchronized input flow, extends the model input to (X_{RGB}, X_A) , where the subscript A signifies the other auxiliary modalities (e.g., depth, thermal infrared, or event). Accordingly, the multi-model tracking can be formulated as $F_{MM} : \{X_{RGB}, X_A, B_0\} \rightarrow B$, where F_{MM} is the multi-modal tracker.

Foundation Model. Typically, the tracker F_{RGB} can be

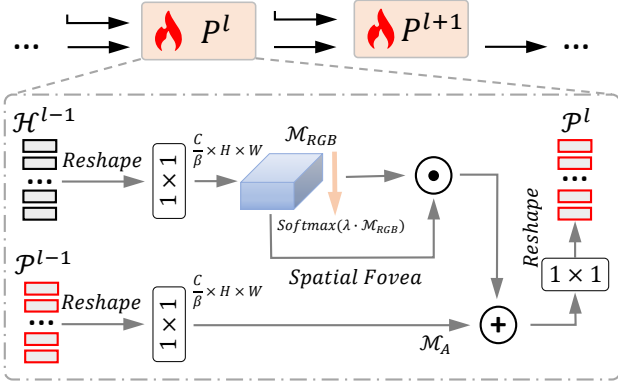


Figure 3. **Detailed design of the proposed MCP.** Multi-modal input flows are sent to the 1×1 convolutions for distilling to the lower-dimension latent space separately. Foundation embeddings first conduct an spatial fovea operation, and then be fused with auxiliary embeddings by element-wise summation. Finally, the multi-modal hybrid embeddings are mapped to the original dimension through a 1×1 convolutional layer.

decomposed into $\phi \circ f$, where $f : \{X_{RGB}, B_0\} \rightarrow \mathcal{H}_{RGB}$ represents the feature extraction and interaction function, and box head $\phi : \mathcal{H}_{RGB} \rightarrow \mathcal{B}$ estimates the final results. Hereby, f is a powerful transformer backbone (*i.e.*, ViT [8]) in our case. Specifically, the input exemplar Z_{RGB} and search frames X_{RGB} are first embedded into patches and flattened to 1D tokens $\mathcal{H}_{RGB}^z \in \mathbb{R}^{N_z \times D}$ and $\mathcal{H}_{RGB}^x \in \mathbb{R}^{N_x \times D}$ with positional embedding [41], where N_z and N_x symbolize the token number of the exemplar and search inputs, and D is the token dimension. Then the token sequences are concatenated to $\mathcal{H}_{RGB}^0 = [\mathcal{H}_{RGB}^z, \mathcal{H}_{RGB}^x]$. We feed the token sequences to an L -layer standard visual transformer encoder. Here we denote \mathcal{H}_{RGB}^{l-1} as inputs to the l -th encoder layer E^l . The forward propagation process is formulated as:

$$\mathcal{H}_{RGB}^l = E^l(\mathcal{H}_{RGB}^{l-1}), \quad l = 1, 2, \dots, L \quad (1)$$

$$\mathcal{B} = \phi(\mathcal{H}_{RGB}^L), \quad (2)$$

where the transformer encoder layer E^l includes Multi-head Self-Attention (MSA), LayerNorm (LN), Feed-Forward Network (FFN), and residual connection, and \mathcal{H}_{RGB}^L is the output of the last encoder layer. The encoder network completes the feature extraction and interaction of exemplar and search patches. We refer readers to OTrack [49] for more details about our RGB-based foundation model.

3.2. Multi-Modal Prompt Tracking

Overall Architecture. Formally, multi-modal tracking provides an extra auxiliary-modal flow which is temporally synchronized and spatially aligned with the RGB flow. As shown in Figure 2, ViPT first feeds the two flows (X_{RGB}

and X_A) to a patch embed layer separately. Each input flow is mapped and flattened into D -dimensional latent space, we denote them as RGB tokens \mathcal{H}_{RGB}^0 and auxiliary tokens \mathcal{H}_A^0 . Then, \mathcal{H}_{RGB}^0 is fed to the foundation model, \mathcal{H}_{RGB}^0 and \mathcal{H}_A^0 are sent into cascade modality-complementary prompter (MCP) to generate modality-specific prompts. The learned prompts are added to the original RGB flow in a form of residuals:

$$\mathcal{H}^l = \mathcal{H}_{RGB}^l + \mathcal{P}^{l+1}, \quad l = 0, 1, \dots, L-1 \quad (3)$$

where \mathcal{H}^l is the prompted tokens that will be injected into the next layer of the foundation model, and \mathcal{P}^{l+1} symbolizes prompt token sequences from the $(l+1)$ -th MCP block. Hereby, we place stage-wise MCP blocks to take full advantage of semantic understanding of diverse levels and diverse modalities. Directly adding prompts to the intermediate features of the foundation model also enables our ViPT to be quickly and easily applied to existing pre-trained foundation trackers. It is worth noting that, different from prompt-tuning approaches (*e.g.*, SPM [29]) which contain trainable prompt-learning network and prediction head, in our ViPT, all the RGB-modal relevant network parameters are frozen, including the patch embedding, feature extraction and interaction, and prediction head.

Modality-Complementary Prompter. Recently, a few works begin to explore introducing some trainable parameters to the frozen pre-trained model to learn effective visual prompts. By fine-tuning only a small number of parameters for prompt learning, they achieve impressive performance on a wide range of vision tasks. A more challenging task, however, is not only to close the gap between upstream and downstream tasks but also to make appropriate and effective harness of inter-modal information. As shown in Figure 2, the proposed MCP blocks are inserted into multiple stages of the foundation network. Formally, providing token sequences $\{\mathcal{H}_{RGB}^0, \mathcal{H}_A^0\}$ of two modalities from a downstream task, and a frozen foundation encoder f_* containing L transformer layers $\{E_*^l\}_{l=1}^L$, the designed MCP is to learn the prompts with these two input flows, the process can be written as:

$$\mathcal{P}^l = P^l(\mathcal{H}^{l-1}, \mathcal{P}^{l-1}), \quad l = 1, 2, \dots, L \quad (4)$$

where P^l denotes the l -th MCP block. Specifically, $\mathcal{H}^0 = \mathcal{H}_{RGB}^0$ and $\mathcal{P}^0 = \mathcal{H}_A^0$. In this way, MCP sufficiently extracts the feature representations of the downstream task at different semantic levels, while learning the complementarity between the two modalities and generating more robust prompts. The blending representation can contribute to a complementarity between the intermediate foundation feature and the learned prompts, and a balance between the foundation and auxiliary modalities.

The detailed design of MCP is shown in Figure 3, our MCP has two input branches for adhibiting the token se-

quences of the foundation flows \mathcal{H}^{l-1} and auxiliary flows \mathcal{P}^{l-1} , respectively. MCP comprises three main steps: (i) projecting to lower-dimensional latent embeddings; (ii) generation of the multi-modal inner complementary representations; and (iii) projecting to original dimension and generating the learned multi-modal visual prompts. Concretely, considering that single-modal flow has certain redundant features, and the prompt block should have as few parameters as possible for practical purposes, we project each flow into $[\frac{C}{\beta} \times H \times W]$ dimension by:

$$\mathcal{M}_{RGB} = g_1(\mathcal{H}^{l-1}), \quad \mathcal{M}_A = g_2(\mathcal{P}^{l-1}), \quad (5)$$

in our case, the input channel C is 768, and the reducing factor β is set to $\frac{768}{8}$ in all the prompt blocks. Hereby, the projection functions $g_1(\cdot)$ and $g_2(\cdot)$ are simple 1×1 convolutional layers. The foundation embeddings \mathcal{M}_{RGB} then perform the spatial fovea operation, which first applies a λ -smoothed spatial *softmax* across all the spatial dimensions, and produces the enhanced embeddings \mathcal{M}_{RGB}^e by applying the channel-wise spatial attention-like mask \mathcal{M}_{fovea} over \mathcal{M}_{RGB} ,

$$\mathcal{M}_{RGB}^e = \mathcal{M}_{RGB} \odot \mathcal{M}_{fovea}, \quad (6)$$

$$\mathcal{M}_{fovea} = \left\{ \frac{e^{\mathcal{M}_{RGB}^{[i,j]}}}{\sum e^{\mathcal{M}_{RGB}^{[i,j]}}} \lambda \right\}, \quad (7)$$

where $i = 1, 2, \dots, H$ and $j = 1, 2, \dots, W$, and λ is a learnable weighted parameter per block. Then, we obtain mixed-modal embeddings by additive binding, and the learned prompts can be gained by,

$$\mathcal{P}^l = g_3(\mathcal{M}_{RGB}^e + \mathcal{M}_A), \quad (8)$$

where $g_3(\cdot)$ is the same as $g_1(\cdot)$ and $g_2(\cdot)$.

Optimization. The multi-modal tracking model F_{MM} is parameter-initialized by the foundation model F_{RGB} . During prompt-tuning, data flows propagate throughout the entire model but we only update the gradient values of a few parameters $\theta = \{\tau^A, \{P^l\}_{l=1}^L\}$ for visual prompt learning, where τ^A denotes the patch embedding functions of auxiliary inputs. Therefore, the default version of our ViPT contains only 0.84M trainable parameters. Despite this, we find that ViPT defeats the full fine-tuning paradigm with two branches for various modalities. Thus, the optimization process can be formulated as

$$\theta_{tuned} = \arg \min_{\theta} \frac{1}{|\mathcal{D}|} \sum \mathcal{L}(\phi(\mathcal{H}^L), \mathcal{B}_{gt}), \quad (9)$$

where \mathcal{D} symbolizes the downstream multi-modal data. By tuning only a few parameters for prompt learning, the model can achieve convergence within a brief period, and effectively leverage the abundant knowledge in the pre-trained foundation model, thereby narrowing the gap between the

models. The overall loss function of ViPT is the same as the foundation model without extra adjusting,

$$\mathcal{L} = L_{cls} + \lambda_{iou} L_{iou} + \lambda_{L_1} L_1. \quad (10)$$

where L_{cls} is the weighted focal loss for classification, L_1 and generalized IoU loss L_{iou} are employed for bounding box regression, λ_{iou} and λ_{L_1} are the regularization parameters, and all the corresponding settings are the same as [49].

3.3. Discussion

The potential of prompt-tuning tracking. We believe that the prompt-tuning paradigm has great potential for multi-modal tracking, mainly in terms of (i) *Prompt-tuning has better adaptability than full fine-tuning, especially for downstream multi-modal tracking where large-scale data is scarce.* Full fine-tuning on the downstream dataset corrupts the quality of pre-trained parameters, so the trackers are more likely to over-fit or reach a sub-optimal state. Especially in the task of RGB+auxiliary modality tracking, the original foundation knowledge can be used directly for the extraction of RGB-modal features. Prompt-tuning fixes the parameters of the foundation model, retaining the knowledge of the pre-trained model, and the prompt module allows for flexible adaptation to task-oriented data. (ii) *Prompt-tuning allows for a closer association between RGB and RGB+auxiliary modality tracking, as well as learning about the modal complementarities.* RGB and auxiliary modalities have different data distributions, so providing an additional feature extraction network for auxiliary-modal inputs may reduce the inter-modal connectivity. (iii) *The remarkable advantages of practical application.* Prompt-tuning tracking has significantly fewer trainable parameters than full fine-tuning, often by an order of magnitude or more. In our work, ViPT contains only 0.84M trainable parameters, less than 1% of the whole model parameters (93.36M). The difference in trainable parameters is further reflected in the more practical aspects of application. Compared with full fine-tuning, prompt-tuning has a shorter training period, at the same time, this paradigm can be deployed to various downstream tracking scenarios without having to store a huge number of foundation parameters repeatedly. These advantages further enables the foundation tracker to be quickly transferred to diverse downstream tracking tasks.

Differences with other prompt-learning methods. The proposed ViPT focuses on effectively introducing prompt-learning mechanisms to multi-modal tracking. In contrast to the existing prompt-learning methods in CV or NLP domains which focus on adaptation between upstream to downstream tasks, ViPT also explores adaptation between single-modal and multi-modal tasks. A similar work to ours is ProTrack [48]. Unfortunately, ProTrack merely introduces the concept of prompt to multi-modal tracking.

	CA3DMS	SiamM_Ds	Siam_LTD	LTDSEd	DAL	CLGS_D	LTMU_B	ATCAIS	DDiMP	DeT	OSTrack	SPT	ProTrack	ViPT
	[31]	[18]	[17]	[18]	[37]	[17]	[17]	[17]	[17]	[47]	[49]	[60]	[48]	(ours)
F-score(↑)	0.223	0.336	0.376	0.405	0.429	0.453	0.460	0.476	0.485	0.532	0.529	0.538	0.578	0.594
Re(↑)	0.228	0.264	0.342	0.382	0.369	0.369	0.417	0.455	0.469	0.506	0.522	0.549	0.573	0.596
Pr(↑)	0.218	0.463	0.418	0.430	0.512	0.584	0.512	0.500	0.503	0.560	0.536	0.527	0.583	0.592

Table 1. Overall performance on DepthTrack test set [47].

	ATOM	DiMP	ATCAIS	DRefine	keep_track	STARK	RGBDDMTracker	DeT	OSTrack	SBT_RGBD	SPT	ProTrack	ViPT
	[7]	[3]	[17]	[19]	[16]	[19]	[16]	[47]	[49]	[16]	[60]	[48]	(ours)
EAO(↑)	0.505	0.543	0.559	0.592	0.606	0.647	0.658	0.657	0.676	0.708	0.651	0.651	0.721
Accuracy(↑)	0.698	0.703	0.761	0.775	0.753	0.803	0.758	0.760	0.803	0.809	0.798	0.801	0.815
Robustness(↑)	0.688	0.731	0.739	0.760	0.797	0.798	0.851	0.845	0.833	0.864	0.851	0.802	0.871

Table 2. Overall performance on VOT-RGBD2022 [16].

Specifically, it conducts a linear summing of multi-modal inputs and uses only RGB-based model parameters without tuning on the downstream data. In contrast, ViPT fully explores the associations between different modalities at different semantic levels, learns the complementarities between modalities in the form of prompt-learning, therefore achieving superior performance.

4. Experiments

4.1. Downstream tasks

ViPT achieves the unification of multiple downstream multi-modal tracking tasks. In this work, three challenging tasks are chosen to validate the effectiveness and generalization of the proposed method. (i) For RGB-D tracking, we evaluate our tracker on Depthtrack [47] and VOT22-RGBD [16]. (ii) For RGB-T tracking, we provide the comparison results on RGBT234 [23] and LasHeR [25]. (iii) For RGB-E tracking, we report the experimental results on the largest VisEvent [43]. We perform prompt-tuning on these downstream tasks without specific modulation, and all the experimental settings are kept to the same.

4.2. Experimental Settings

When fine-tuning our ViPT on downstream tasks, we choose train sets of Depthtrack [47] for RGB-D tracking, LasHeR [25] for RGB-T tracking, and VisEvent [43] for RGB-E tracking. Some recent works (*e.g.*, DeT [47] and mfDiMP [53]) employ other algorithms to generate more sample pairs for downstream tasks, thus alleviating the problem of insufficient data in downstream tasks. While in our work, to demonstrate the effectiveness of prompt-tuning, we only use the most basic multi-modal training set mentioned above. ViPT is trained on 2 NVIDIA Tesla A100 GPUs with a global batch size of 64. The model fine-tuning takes 60 epochs that each epoch contains 6×10^4 sample pairs. The AdamW optimizer [32] with a weight decay of 10^{-4} is adopted. The initial learning rate is set to 4×10^{-5} and decreased by the factor of 10 after 48 epochs. The

fixed parameters are initialized with the pre-trained foundation model [49] in RGB-based tracking, and the trainable prompt learning parameters are initialized with xavier uniform initialization scheme [11].

4.3. Main Properties and Analysis

DepthTrack. DepthTrack [47] is a large-scale long-term RGB-D tracking benchmark, which contains 150 training and 50 testing videos with 15 per-frame attributes. It uses precision (Pr) and recall (Re) to measure the accuracy and robustness of target localization. F-score, calculated by $F = \frac{2RePr}{Re+Pr}$, is the primary measure. Though our ViPT is a short-term algorithm, Table 1 shows ViPT surpasses all previous SOTA trackers and obtains the highest F-score of 59.4%, gaining a significant improvement of 6.5% compared with our foundation model [49].

VOT-RGBD2022. VOT-RGBD2022 [16] is an up-to-date benchmark that contains 127 short-term RGB-D sequences for exploiting the role of depth in RGB-D tracking. It applies an anchor-based short-term evaluation protocol introduced in [17] that needs trackers to multi-start from different initialization points. The expected average overlap (EAO) is the overall performance measure. As reported in Table 2, our ViPT is superior to previous competitive methods, getting an EAO of 0.721. ViPT outperforms the foundation model by 4.5% EAO while another prompt-




Figure 4. Overall performance on LasHeR test set [25].

	mfDiMP	SGT	DAFNet	OSTrack	FANet	MaCNet	CAT	JMMAC	CMPP	APFNet	ProTrack	ViPT
	[53]	[26]	[10]	[49]	[62]	[50]	[24]	[55]	[42]	[44]	[48]	(ours)
MPR(\uparrow)	0.646	0.720	0.796	0.729	0.787	0.790	0.804	0.790	0.823	0.827	0.795	0.835
MSR(\uparrow)	0.428	0.472	0.544	0.549	0.553	0.554	0.561	0.573	0.575	0.579	0.599	0.617

Table 3. Overall performance on RGBT234 dataset [23].

based tracker ProTrack [48], only has an improvement of 0.4% relative to their baseline STARK_RGBD [46].

RGBT234. RGBT234 [23] is a large-scale RGB-T tracking benchmark, which contains 234 videos with visible and thermal infrared pairs. MSR and MPR are adopted for performance evaluation. Here we compare our proposed ViPT with recent RGB-T trackers. As shown in Table 3, our ViPT reaches the top MSR and MPR of 61.7% and 83.5%, respectively, which beats the well-designed RGB-T trackers and outperforms ProTrack [48] by 1.8% on MSR.

LasHeR. LasHeR [25] is a large-scale high-diversity benchmark for short-term RGB-T tracking. We evaluate the trackers on 245 testing video sequences in terms of precision plot and success plot. The results are reported in Figure 4. Surprisingly, ViPT surpasses previous SOTA algorithms by a large margin, which exceeds the second place by 10.5% and 11.3% on success and precision, respectively. It can be inferred that our prompt-tuning paradigm effectively utilizes the thermal information and makes ViPT adapt well to complex tracking scenarios.

VisEvent. VisEvent [43] is the largest visible-event benchmark dataset collected from real-world for single object tracking at present and we perform comparisons on its 320 testing videos. Note we just use event images transformed from raw event data rather than event flow. Figure 5 shows ViPT achieves a gain of 5.8% on success and 6.3% on precision over the runner-up OSTrack [49] and outperforms other visible-event SOTA trackers.

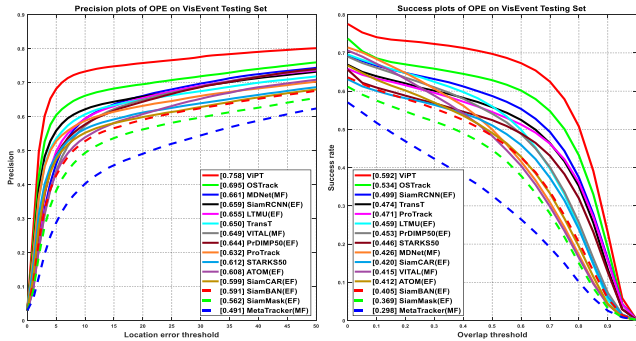


Figure 5. Overall performance on VisEvent test set [43].

4.4. Exploration Studies

We further explore the characteristics of our ViPT on multiple downstream tasks of multiple representative benchmarks. We present the experimental results on Depth-

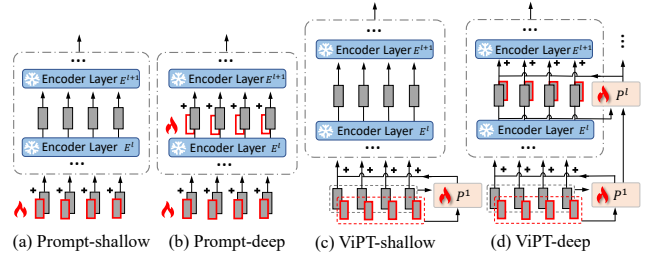


Figure 6. Variants of vanilla prompt-structure and ViPT.

track [47], LasHeR [25], and VisEvent [43].

Power of Multiple Modalities. To investigate the benefits gained from multi-modal flows, we evaluate the performance with single-modal inputs (RGB-based foundation tracker) and multi-modal inputs. For multi-modal cases, we further compare the effectiveness between FFT and our ViPT. Specifically, FFT is the full fine-tuning tracker which is extended from the foundation model by adding an auxiliary-modal branch, and the branch design and modalities interaction are inspired by [47]. As shown in Table 4, after combining the auxiliary-modal inputs, FFT shows remarkable improvement over the foundation model, especially on LasHeR (10.5% \uparrow). It is encouraging to note that ViPT still beats FFT on performance while having two orders of magnitude fewer trainable parameters.

Method	Params \uparrow	Depthtrack [47]			LasHeR [25]		VisEvent [43]	
		F-score(\uparrow)	Re(\uparrow)	Pr(\uparrow)	SR(\uparrow)	PR(\uparrow)	SR(\uparrow)	PR(\uparrow)
Foundation	-	52.9	52.2	53.6	41.2	51.5	53.4	69.5
FFT	178.6M (100%)	55.6 (2.7 \uparrow)	55.4	55.7	51.7 (10.5 \uparrow)	64.8	58.7 (5.3 \uparrow)	75.2
ViPT	0.84M (0.9%)	59.4 (6.5\uparrow)	59.6	59.2	52.5 (11.3\uparrow)	65.1	59.2 (5.8\uparrow)	75.8
Prompt-shaw	0.59M (0.6%)	52.0	51.5	52.5	47.2	58.7	53.7	70.2
Prompt-deep	3.29M (3.4%)	54.9	54.7	55.1	50.7	63.1	54.8	71.4
ViPT-shaw	0.61M (0.7%)	56.4	56.7	56.2	47.9	59.6	57.8	74.9

Table 4. Ablation studies on multiple downstream benchmarks. Params \uparrow denotes the number of trainable parameters, and its proportion of the overall model parameters is noted in parentheses. Results are reported in percentage (%).

Variants Comparison. We further explore different variant versions of our ViPT. As shown in Figure 6 (c) & (d), we provide two variants of ViPT, ViPT-shallow (ViPT-shaw in Table 4) and ViPT-deep (ViPT in Table 4). Moreover, to verify the validity of our MCP module, we introduce prompt tokens into the foundation model with VPT-style [14] (Figure 6 (a) & (b)). Note that for receiving auxiliary-modal inputs, we replace the input prompts with the embeddings of auxiliary modality. And we choose to

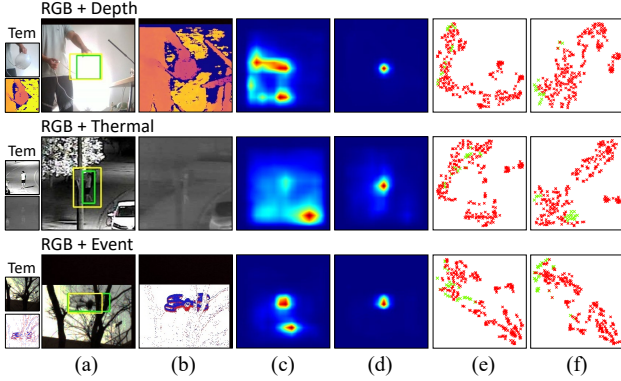


Figure 7. Visualization of response and t-SNE [40] maps. (a) RGB flows. The yellow and green bounding boxes denote Foundation tracker and ViPT, respectively. (b) Auxiliary-modal flows. (c) Score maps of foundation model. (d) Score maps of ViPT. (e) t-SNE maps of the foundation model. (f) t-SNE maps of ViPT.

introduce the prompt parameters as a summation instead of the concatenation to align RGB and other auxiliary modalities. The trainable parameters and performance comparison are reported in Table 4. We can see that ViPT achieves favorable precision-efficiency trade-offs.

Number Analysis of Prompt Block. ViPT achieves multi-modal tracking by inserting MCP blocks into different semantic layers of the foundation model. Here we experimentally investigate the effect of providing different numbers of prompts to the foundation model. Specifically, we set different placement intervals for the inserted blocks, and we insert MCP per 1, 2, 4, 6, and 12 layers of the foundation encoder. Note that the first and the last case are the variants of our ViPT-deep and ViPT-shallow. As shown in Figure 8, we observed that as the number of inserted MCP blocks increases, there is a corresponding increase in performance on all benchmarks.




Figure 8. Influence of the number of MCP blocks. The number of trainable parameters is shown in each subfigure for reference.

Training Volume and Tunable Parameters. We note that DeT [47] and mfDiMP [53] expand the training data by generating more RGB-D and RGB-T image pairs. Hence, we further expand the training data to see its impact. In the implementation, we use the same data generation method as them. Moreover, we unfreeze the overall parameters of ViPT to see the impact of number of trainable parameters. The results are reported in Table 5. We can see that only expanding the train set or unfreezing the rest param-

eters can not bring better results. The reason may be that a small number of trainable parameters is more conducive for model learning in downstream tasks with limited train data. Conducting them together may gain some benefits (*e.g.*, on LasHeR), but it will be more costly compared with ViPT.

Model	ET	UFP	Depthtrack [47]			LasHeR [25]	
			F-score(↑)	Re(↑)	Pr(↑)	SR(↑)	PR(↑)
ViPT			59.4	59.6	59.2	52.5	65.1
①	✓		59.3	59.3	59.2	49.7	61.6
②		✓	55.4	55.3	55.6	51.3	64.5
③	✓	✓	58.3	58.6	58.0	53.9	67.0

Table 5. Ablation studies on training volume and tunable parameters. ET: expanding the training data. UFP: unfreeze all the parameters tunably. Results are reported in percentage (%).

Visualization. To explore how the learned modality-complementary prompts work in ViPT, we visualize several representative response maps and the corresponding t-SNE maps in Figure 7. As can be seen, with auxiliary-modal prompts, ViPT has more unambiguous and discriminative responses when facing some complex scenarios, such as illumination variation and background cluster. t-SNE maps also verify that learned prompts help the foundation tracker distinguish the objects from the background due to the more separable distribution boundaries. As a result, it leads to more accurate and robust object localization overall.

5. Conclusion

In this work, we present ViPT, a new parameter-efficient vision tuning framework by introducing prompt-learning ideology to multi-modal tracking. The core idea of ViPT is to maximally exploit the pre-trained foundation model which is trained at scale, and excavate the relevance of foundation and downstream tracking, and the complementarity of multiple modalities. Extensive experiments on multiple downstream tasks demonstrate the effectiveness and generalization of ViPT. We expect this work can attract more attention to prompt learning for multi-modal tracking.

Limitations. Despite achieving a unified architecture for multi-modal tracking tasks, ViPT mainly focuses on visual prompting. However, some vision-language tasks are emerging. ViPT may have the potential to be extended to more tracking tasks like vision-language tracking. Moreover, our method still needs to train separately for the different multi-modal tasks. In the future, we are interested in joint training a general model for multiple modalities.

Acknowledgements. This work was supported in part by the National Key R&D Program of China under Grant No. 2018AAA0102001, the National Natural Science Foundation of China under Grant Nos. 62293542, U1903215, and the Fundamental Research Funds for the Central Universities under Grant No. DUT22ZD210.

References

- [1] Hyojin Bahng, Ali Jahanian, Swami Sankaranarayanan, and Phillip Isola. Exploring visual prompts for adapting large-scale models. *arXiv preprint arXiv:2203.17274*, 1(3):4, 2022. 2
- [2] Luca Bertinetto, Jack Valmadre, João F Henriques, Andrea Vedaldi, and Philip H S Torr. Fully-convolutional siamese networks for object tracking. In *ECCVW*, 2016. 1, 2
- [3] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *ICCV*, 2019. 1, 2, 6
- [4] Shoufa Chen, Chongjian Ge, Zhan Tong, Jiangliu Wang, Yibing Song, Jue Wang, and Ping Luo. Adaptformer: Adapting vision transformers for scalable visual recognition. *NeurIPS*, 2022. 3
- [5] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *CVPR*, 2021. 1, 2
- [6] Ciarán Ó Conaire, Noel E O’Connor, and Alan Smeaton. Thermo-visual feature fusion for object tracking using multiple spatiogram trackers. *Machine Vision and Applications*, 19(5):483–494, 2008. 2
- [7] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. ATOM: Accurate tracking by overlap maximization. In *CVPR*, 2019. 1, 6
- [8] 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. *ICLR*, 2021. 4
- [9] 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. 1, 2
- [10] Yuan Gao, Chenglong Li, Yabin Zhu, Jin Tang, Tao He, and Futian Wang. Deep adaptive fusion network for high performance RGBT tracking. In *ICCVW*, pages 0–0, 2019. 7
- [11] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *ICAI*, 2010. 6
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2
- [13] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *TPAMI*, 2019. 1, 2
- [14] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, 2022. 2, 3, 7
- [15] Shibo Jie and Zhi-Hong Deng. Convolutional bypasses are better vision transformer adapters. *arXiv preprint arXiv:2207.07039*, 2022. 3
- [16] 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*, pages 431–460. Springer, 2023. 6
- [17] 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*, pages 547–601. Springer, 2020. 6
- [18] Matej Kristan, Jiri Matas, Ales Leonardis, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kamarainen, Luka Čehovin Zajc, Ondrej Drbohlav, Alan Lukežic, Amanda Berg, et al. The seventh visual object tracking vot2019 challenge results. In *ICCVW*, pages 0–0, 2019. 6
- [19] 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*, pages 2711–2738, 2021. 6
- [20] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *EMNLP*, 2021. 3
- [21] 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. 1
- [22] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *CVPR*, 2018. 1, 2
- [23] Chenglong Li, Xinyan Liang, Yijuan Lu, Nan Zhao, and Jin Tang. RGB-T object tracking: Benchmark and baseline. *Pattern Recognition*, 96:106977, 2019. 6, 7
- [24] Chenglong Li, Lei Liu, Andong Lu, Qing Ji, and Jin Tang. Challenge-aware RGBT tracking. In *ECCV*, pages 222–237. Springer, 2020. 7
- [25] 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 Transactions on Image Processing*, 31:392–404, 2021. 1, 6, 7, 8
- [26] Chenglong Li, Nan Zhao, Yijuan Lu, Chengli Zhu, and Jin Tang. Weighted sparse representation regularized graph learning for RGB-T object tracking. In *ACMMM*, pages 1856–1864, 2017. 7
- [27] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 1, 3
- [28] Hongjie Liu, Diederik Paul Moeyss, Gautham Das, Daniel Neil, Shih-Chii Liu, and Tobi Delbrück. Combined frame- and event-based detection and tracking. In *ISCV*, pages 2511–2514. IEEE, 2016. 2
- [29] Lingbo Liu, Bruce XB Yu, Jianlong Chang, Qi Tian, and Chang-Wen Chen. Prompt-matched semantic segmentation. *arXiv preprint arXiv:2208.10159*, 2022. 4
- [30] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*, 2021. 3

- [31] Ye Liu, Xiao-Yuan Jing, Jianhui Nie, Hao Gao, Jun Liu, and Guo-Ping Jiang. Context-aware three-dimensional mean-shift with occlusion handling for robust object tracking in RGB-D videos. *IEEE Transactions on Multimedia*, 21(3):664–677, 2018. 2, 6
- [32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. 6
- [33] Chengwei Luo, Bin Sun, Ke Yang, Taoran Lu, and Wei-Chang Yeh. Thermal infrared and visible sequences fusion tracking based on a hybrid tracking framework with adaptive weighting scheme. *Infrared Physics & Technology*, 99:265–276, 2019. 2
- [34] S Mahdi H Miangoleh, Sebastian Dille, Long Mai, Sylvain Paris, and Yagiz Aksoy. Boosting monocular depth estimation models to high-resolution via content-adaptive multi-resolution merging. In *CVPR*, pages 9685–9694, 2021. 3
- [35] Franziska Mueller, Dushyant Mehta, Oleksandr Sotnychenko, Srinath Sridhar, Dan Casas, and Christian Theobalt. Real-time hand tracking under occlusion from an egocentric RGB-D sensor. In *ICCV*, pages 1154–1163, 2017. 2
- [36] 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. 1, 2
- [37] Yanlin Qian, Song Yan, Alan Lukežič, Matej Kristan, Joni Kristian Kämäräinen, and Jiří Matas. Dal: A deep depth-aware long-term tracker. In *ICPR*, pages 7825–7832. IEEE, 2021. 6
- [38] 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. PMLR, 2021. 2
- [39] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, and Michael Bernstein. ImageNet Large scale visual recognition challenge. *IJCV*, 2015. 1, 2
- [40] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008. 8
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. 2, 4
- [42] 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. 7
- [43] 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. 1, 2, 6, 7
- [44] Yun Xiao, Mengmeng Yang, Chenglong Li, Lei Liu, and Jin Tang. Attribute-based progressive fusion network for RGBT tracking. In *AAAI*, 2022. 1, 7
- [45] Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *ECCV*, pages 585–601, 2018. 1
- [46] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *ICCV*, 2021. 7
- [47] Song Yan, Jinyu Yang, Jani Käpylä, Feng Zheng, Aleš Leonardis, and Joni-Kristian Kämäräinen. Depthtrack: Unveiling the power of RGBD tracking. In *ICCV*, pages 10725–10733, 2021. 1, 2, 6, 7, 8
- [48] Jinyu Yang, Zhe Li, Feng Zheng, Ales Leonardis, and Jingkuan Song. Prompting for multi-modal tracking. In *ACMMM*, pages 3492–3500, 2022. 3, 5, 6, 7
- [49] 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. Springer, 2022. 4, 5, 6, 7
- [50] 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. 7
- [51] 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*, pages 8801–8810, 2022. 1
- [52] Jiqing Zhang, Xin Yang, Yingkai Fu, Xiaopeng Wei, Baocai Yin, and Bo Dong. Object tracking by jointly exploiting frame and event domain. In *ICCV*, pages 13043–13052, 2021. 1, 2
- [53] 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*, pages 0–0, 2019. 2, 6, 7, 8
- [54] Pengyu Zhang, Dong Wang, Huchuan Lu, and Xiaoyun Yang. Learning adaptive attribute-driven representation for real-time RGB-T tracking. *IJCV*, 129(9):2714–2729, 2021. 2
- [55] Pengyu Zhang, Jie Zhao, Chunjuan Bo, Dong Wang, Huchuan Lu, and Xiaoyun Yang. Jointly modeling motion and appearance cues for robust RGB-T tracking. *IEEE Transactions on Image Processing*, 30:3335–3347, 2021. 2, 7
- [56] Pengyu Zhang, Jie Zhao, Dong Wang, Huchuan Lu, and Xiang Ruan. Visible-thermal uav tracking: A large-scale benchmark and new baseline. In *CVPR*, pages 8886–8895, 2022. 1
- [57] Tianlu Zhang, Xueru Liu, Qiang Zhang, and Jungong Han. Siamcda: Complementarity-and distractor-aware RGB-T tracking based on siamese network. *TCSVT*, 32(3):1403–1417, 2021. 1
- [58] Xingchen Zhang, Ping Ye, Dan Qiao, Junhao Zhao, Shengyun Peng, and Gang Xiao. Object fusion tracking based on visible and infrared images using fully convolutional siamese networks. In *FUSION*, pages 1–8. IEEE, 2019. 2
- [59] Zangwei Zheng, Xiangyu Yue, Kai Wang, and Yang You. Prompt vision transformer for domain generalization. *arXiv preprint arXiv:2208.08914*, 2022. 2

- [60] 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. [1](#), [2](#), [6](#)
- [61] Yabin Zhu, Chenglong Li, Bin Luo, Jin Tang, and Xiao Wang. Dense feature aggregation and pruning for RGBT tracking. In *ACMMM*, pages 465–472, 2019. [1](#), [2](#)
- [62] Yabin Zhu, Chenglong Li, Jin Tang, and Bin Luo. Quality-aware feature aggregation network for robust RGBT tracking. *IEEE Transactions on Intelligent Vehicles*, 6(1):121–130, 2020. [7](#)