

# Group Orthogonal Low-Rank Adaptation for RGB-T Tracking

Zekai Shao, Yufan Hu, Jingyuan Liu, Bin Fan, Hongmin Liu\*

School of Intelligence Science and Technology, University of Science and Technology Beijing  
 melan.shao@foxmail.com, huyufanqiaxuan@gmail.com, lly.ustb@foxmail.com,  
 bin.fan@ieee.org, hmliu\_82@163.com

## Abstract

Parameter-efficient fine-tuning has emerged as a promising paradigm in RGB-T tracking, enabling downstream task adaptation by freezing pretrained parameters and fine-tuning only a small set of parameters. This set forms a rank space made up of multiple individual ranks, whose expressiveness directly shapes the model’s adaptability. However, quantitative analysis reveals low-rank adaptation exhibits significant redundancy in the rank space, with many ranks contributing almost no practical information. This hinders the model’s ability to learn more diverse knowledge to address the various challenges in RGB-T tracking. To address this issue, we propose the Group Orthogonal Low-Rank Adaptation (GOLA) framework for RGB-T tracking, which effectively leverages the rank space through structured parameter learning. Specifically, we adopt a rank decomposition partitioning strategy utilizing singular value decomposition to quantify rank importance, freeze crucial ranks to preserve the pretrained priors, and cluster the redundant ranks into groups to prepare for subsequent orthogonal constraints. We further design an inter-group orthogonal constraint strategy. This constraint enforces orthogonality between rank groups, compelling them to learn complementary features that target diverse challenges, thereby alleviating information redundancy. Experimental results demonstrate that GOLA effectively reduces parameter redundancy and enhances feature representation capabilities, significantly outperforming state-of-the-art methods across four benchmark datasets and validating its effectiveness in RGB-T tracking tasks.

**Code** — <https://github.com/MelanTech/GOLA>

## Introduction

RGB-T tracking combines the complementary strengths of visible and infrared data, significantly enhancing the robustness in challenging scenarios, such as low-light or occluded conditions. It has been applied in multiple domains, demonstrating its effectiveness in complex scenarios.

Recently, most RGB-T trackers adopt a prompt-tuning paradigm, integrating lightweight prompt modules (Zhu

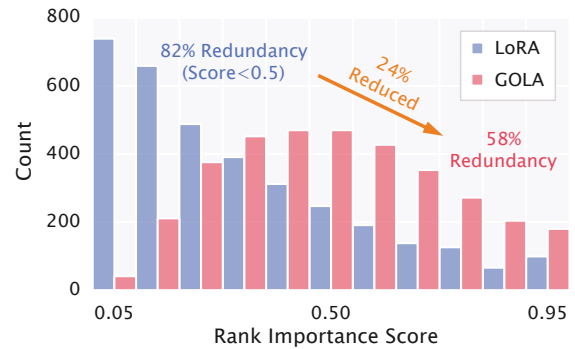


Figure 1: Comparison of rank importance score distribution between LoRA and our proposed GOLA. The rank space of LoRA exhibits significant redundancy, while our proposed GOLA effectively mitigates this issue.

et al. 2023; Cao et al. 2024; Hu et al. 2025b) to facilitate multimodal fusion while maintaining an acceptable speed. However, this prompt-tuning approach typically relies on feature dimensionality reduction to balance efficiency and performance, which limits the capacity of interaction between modalities to express information in a low-rank space and weakens the model’s overall performance. LoRA (Hu et al. 2022) applies the low-rank concept to the parameter level by freezing the backbone and optimizing only the low-rank matrices, thereby avoiding information loss and enabling parameter merging during the inference stage, which enhances inference efficiency. Nonetheless, in complex RGB-T tracking tasks, the redundant ranks of LoRA do not effectively facilitate the acquisition of diverse knowledge, which directly undermines its ability to address various challenges. As shown in Fig. 1, after training the model with LoRA, we decompose all ranks in its parameters and plot the histogram of singular values. The results indicate significant redundancy in the rank space of LoRA, with only a few ranks becoming dominant. This phenomenon arises from LoRA’s compression of all information into a low-rank space during downstream adaptation, making it difficult to address the diverse challenges present in RGB-T tracking. In training, to achieve rapid convergence to an optimal solution, it tends to prioritize the integration of a few crucial ranks that en-

\*Corresponding author.

capsulate the most significant task, which overshadows the pretrained priors. Simultaneously, the remaining ranks remain unactivated due to a lack of targeted optimization signals, failing to capture the fine-grained features necessary for specific challenges, while also struggling to supplement the expressive blind spots uncovered by crucial ranks. Ultimately, this severely limits the overall expressiveness of the rank space, hampering the model’s ability to handle diverse challenges in RGB-T tracking.

To address the issues above, we propose a novel Group Orthogonal Low-Rank Adaptation (GOLA) framework for RGB-T tracking, aiming to preserve the pretrained priors while fully activating the expressive potential of the redundant rank space. First, we introduce a rank decomposition partitioning strategy that quantifies the significance of each rank by performing singular value decomposition on the LoRA parameter matrix and categorizes these ranks into crucial and redundant ones. We freeze the crucial ranks to retain the strong generalization capability of the pretrained weights. This operation is conducted offline and does not incur additional training burdens. Furthermore, we utilize clustering to partition the redundant rank components into parallel rank subsets, establishing a foundation for subsequent orthogonal constraints. We then introduce an inter-group orthogonal constraint strategy by constructing an orthogonal loss, which enforces the parameters of two randomly sampled groups to maintain an orthogonal relationship. This strategy ensures that the parameters of each group learn mutually independent and complementary feature transformations, effectively avoiding overlap and redundancy in feature learning. Ultimately, the frozen critical ranks robustly preserve the existing priors, while the grouped redundant ranks focus on learning knowledge tailored to different challenges under orthogonal constraints. This synergistic effect collectively enhances the model’s tracking performance in complex RGB-T scenarios. In summary, our contributions are threefold:

- We propose the Group Orthogonal Low-Rank Adaptation (GOLA) framework for RGB-T tracking, addressing the shortcomings of rank space redundancy and limited expressive capability, enhancing the model’s ability to handle diverse challenges.
- In our framework, we design a rank decomposition strategy to freeze the crucial ranks and group the redundant ranks, along with an inter-group orthogonal constraint that enforces complementary feature learning across groups to enhance the model’s representation ability.
- Experimental results validate the effectiveness of the proposed method in RGB-T tracking, demonstrating that our method achieves optimal performance and efficiency.

## Related Works

### RGB-T Tracking

With the aid of visible and infrared modalities, RGB-T tracking significantly enhances tracking robustness in complex scenes through information fusion. Recent RGB-T

trackers are mainly categorized into full fine-tuning, prompt-tuning, online adaptation, and low-rank adaptation methods based on their training strategies. Full fine-tuning approaches fine-tune trackers pretrained on the RGB modality to adapt to the RGB-T task. TBSI (Hui et al. 2023) uses a dual-stream backbone, treats templates as a medium to bridge RGB and infrared modalities, and incorporates multi-modal information via template updates. USTrack (Xia et al. 2024) first unifies features into a single backbone, solving issues of separate modal feature extraction and slow speed. Full fine-tuning achieves good tracking performance but has long training cycles and high resource consumption. Moreover, scarce RGB-T training data makes existing methods prone to overfitting and performance degradation during full fine-tuning. Prompt-tuning-based methods (Zhu et al. 2023; Wang et al. 2024; Zeng et al. 2024) have reduced training overhead by fine-tuning lightweight prompt modules, while maintaining tracking performance comparable to that of full fine-tuning methods. Nevertheless, prompt-tuning-based methods rely on additional modules and parameters, severely limiting model running efficiency. PURA (Shao et al. 2025) proposes a framework based on parameter decomposition and recovery to achieve adaptation for RGB-T tracking in domain shift scenarios. UnTrack (Wu et al. 2024) introduces LoRA (Hu et al. 2022) for low-rank fine-tuning of the backbone to reduce the number of trainable parameters. However, it still employs prompt-tuning methods for the infrared modality branch, which similarly relies on additional parameters. Moreover, the low-rank characteristics of LoRA hinder its ability to learn diverse knowledge to tackle different challenges. In this work, we focus on reducing parameter redundancy through the structured management of rank spaces and orthogonal constraints, thereby achieving efficient cross-modal transfer.

### Parameter-Efficient Fine-Tuning

As the core paradigm for lightweight adaptation of pretrained models, Parameter-Efficient Fine-Tuning (PEFT) has recently demonstrated successful extensions to computer vision, providing solutions for model training in resource-constrained scenarios. PEFT can be broadly categorized into adapter-based methods, prompt-based methods, and low-rank adaptation (LoRA) (Hu et al. 2022). Adapter-based methods (Houlsby et al. 2019; Jie et al. 2024) enable transfer learning by introducing lightweight modules into pretrained models. Prompt-based methods (Jia et al. 2022; Li and Liang 2021; Lester, Al-Rfou, and Constant 2021) focus on utilizing carefully designed prompts to guide pretrained models in generating responses for specific tasks. LoRA achieves fine-tuning by decomposing weight updates into low-rank matrices during the adaptation process, leveraging the inherent low-rank structure of weight updates to enable adaptation with low computational overhead. To enhance learning efficiency, some variants (Zhang et al. 2023; Liu et al. 2024b) dynamically adjust parameters through weight decomposition, optimizing the learning process. Recently, methods based on a mixture of experts have shown significant effectiveness in multitask scenarios. MOLA (Gao et al. 2024) markedly improves fine-tuning outcomes by setting

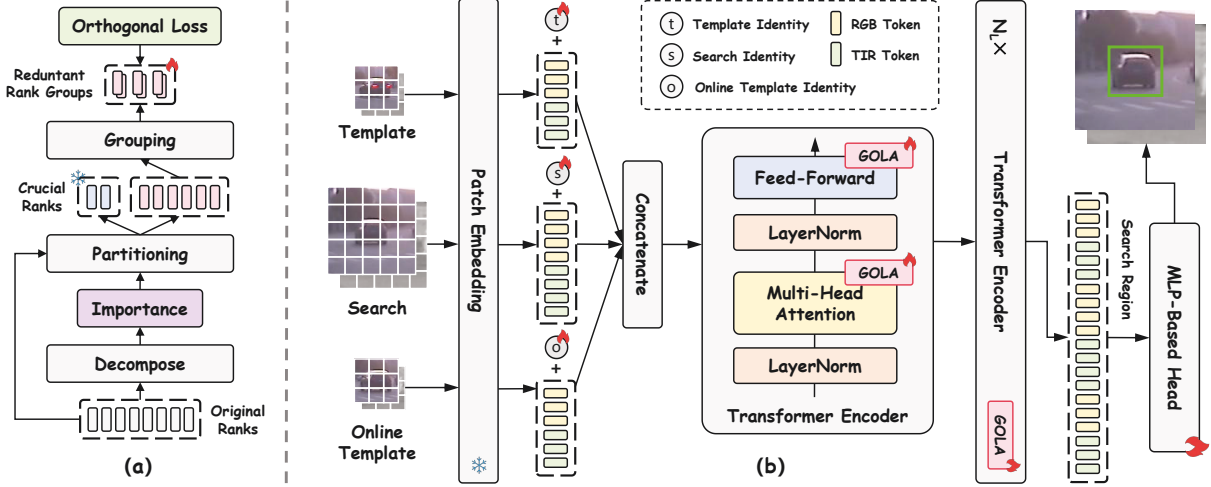


Figure 2: (a) Our proposed Group Orthogonal Low-Rank Adaptation (GOLA) framework. We decompose pretrained ranks into crucial ranks and redundant rank groups, freezing the crucial ranks to retain generalization. Our inter-group orthogonal constraint boosts the expressiveness of the redundant rank groups, fully utilizing the redundant rank space. (b) The overall architecture of our tracking framework. GOLA is applied to each linear layer of the backbone for fine-tuning.

varying numbers of experts at different layers. MixLoRA (Li et al. 2024) employs multiple LoRA-based experts while introducing load-balancing losses to address the imbalance issues among routers. Furthermore, LoRA has revealed significant potential in tracking tasks, with LoRAT (Lin et al. 2024) and UnTrack (Wu et al. 2024) utilizing this technique for the fine-tuning of the backbone, thereby enhancing tracking performance while reducing the computational burden.

## Methodology

This paper presents the Group Orthogonal Low-Rank Adaptation (GOLA) framework for RGB-T tracking, which enhances tracking performance by rank partition and orthogonal constraints on pretrained ranks. This section outlines the tracker architecture and provides details on the method’s implementation. The overall framework is depicted in Fig. 2.

### Tracking Framework

As shown in Fig. 2, our method adopts a simple single-stream tracking framework. Let  $\mathbf{I}_v^z, \mathbf{I}_t^z$  be template images,  $\mathbf{I}_v^x, \mathbf{I}_t^x$  be search images, and  $\mathbf{I}_v^o, \mathbf{I}_t^o$  be online template images. Here,  $v$  and  $t$  represent the visible and thermal modalities, respectively. We first transform these images into token sequences, resulting in  $\mathbf{Z}_v, \mathbf{Z}_t \in \mathbb{R}^{n_z \times c}$ ,  $\mathbf{X}_v, \mathbf{X}_t \in \mathbb{R}^{n_x \times c}$ , and  $\mathbf{O}_v, \mathbf{O}_t \in \mathbb{R}^{n_o \times c}$ . Here,  $\mathbf{Z}$ ,  $\mathbf{X}$ , and  $\mathbf{O}$  represent the template, the search region, and the online template, respectively.  $n_z$  and  $n_x$  represent the number of tokens for the template and the search region, respectively, and  $c$  signifies the number of feature channels. In addition, we provide unique trainable type embeddings (Lin et al. 2024) for each category of tokens to prevent information confusion.

The tokens are then concatenated to form the final embedding vector denoted as  $h = [\mathbf{Z}_v; \mathbf{Z}_t; \mathbf{X}_v; \mathbf{X}_t; \mathbf{O}_v; \mathbf{O}_t]$ . Subsequently, we process  $h$  through a series of encoder layers for joint feature extraction. Our GOLA uses the same

fine-tuning strategy as LoRA (Hu et al. 2022). For each linear layer parameter  $\mathbf{W}$  in the backbone, we introduce two low-rank matrices  $\mathbf{A}$  and  $\mathbf{B}$  to update it as follows:

$$h' = \mathbf{W}h + \mathbf{B}\mathbf{A}h, \quad (1)$$

where  $\mathbf{A} = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_r\} \in \mathbb{R}^{r \times c}$  and  $\mathbf{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_r\} \in \mathbb{R}^{c \times r}$ , and  $r \ll c$ . In the inference stage, we use the following formula for parameter merging:

$$\mathbf{W}' = \mathbf{W} + \mathbf{B}\mathbf{A}. \quad (2)$$

Subsequently, the search region features produced by the backbone are fed into the prediction head to obtain the final results. Additionally, we use the max score map value from the prediction head as confidence, comparing it with the threshold  $\tau$  to update the online template.

### Group Orthogonal Low-Rank Adaptation

**Rank Decomposition Partition.** Our rank decomposition partition aims to leverage singular value decomposition to quantify the importance of ranks, thereby distinguishing between critical and redundant ranks for structured learning. Notably, this process is offline and does not affect the training workflow. The matrix  $\mathbf{B}$  exhibits a significant rank correlation with task specificity compared to the matrix  $\mathbf{A}$ , which functions as a general feature extractor (Tian et al. 2024). This stronger association of  $\mathbf{B}$  with modality-specific tasks leads to its use as the basis for partitioning. Specifically, we first perform singular value decomposition on matrix  $\mathbf{B}$  to obtain the rank importance:

$$\Sigma, \mathbf{V} \leftarrow \text{SVD}(\bar{\mathbf{B}}), \quad (3)$$

where  $\bar{\mathbf{B}}$  represents the mean-centered matrix of  $\mathbf{B}$ , singular values  $\Sigma \in \mathbb{R}^r$  represent the importance of each rank, and the singular vectors  $\mathbf{V} \in \mathbb{R}^{c \times c}$  represent the principal direction of each rank, with both ordered in descending order

of importance. Subsequently, we select the top- $k$  singular vectors  $V_k \in \mathbb{R}^{k \times c}$  as reference vectors and use the top- $k$  singular values  $\Sigma_k \in \mathbb{R}^k$  as weights. We compute the importance score  $\mathbf{S}$  for each rank by calculating the weighted similarity between each rank and the reference ranks. This process is represented as:

$$\mathbf{S} = \|\bar{\mathbf{B}}^\top \mathbf{V}_k^\top \odot \Sigma_k\|_2, \quad (4)$$

where  $\odot$  represents the hadamard product, and  $\|\cdot\|_2$  denotes  $L_2$  column normalization. After obtaining the importance score  $\mathbf{S} \in \mathbb{R}^r$ , we sort all ranks in descending order based on  $\mathbf{S}$  to determine their priority. Let the sorted index be denoted as  $\sigma$ , which satisfies:

$$\mathbf{S}_{\sigma_1} \geq \mathbf{S}_{\sigma_2} \geq \dots \geq \mathbf{S}_{\sigma_r}, \quad (5)$$

where  $\sigma_i$  denotes the original index of the  $i$ -th sorted element after sorting. Subsequently, we perform the same rank partitioning on  $\mathbf{A}$  and  $\mathbf{B}$  to obtain crucial ranks  $\mathbf{A}_c$  and  $\mathbf{B}_c$ :

$$\begin{aligned} \mathbf{A}_c &= \{\mathbf{a}_{\sigma_1}, \mathbf{a}_{\sigma_2}, \dots, \mathbf{a}_{\sigma_k}\}, \\ \mathbf{B}_c &= \{\mathbf{b}_{\sigma_1}, \mathbf{b}_{\sigma_2}, \dots, \mathbf{b}_{\sigma_k}\}, \end{aligned} \quad (6)$$

where  $\mathbf{a}_{\sigma_i}$  and  $\mathbf{b}_{\sigma_i}$  represent the selections of ranks using the  $i$ -th sorted index. It is noteworthy that performing the same sorting operation on  $\mathbf{A}$  and  $\mathbf{B}$  produces new weights that are equivalent to the original weights, thereby not affecting the training process. Since  $\mathbf{A}_c$  and  $\mathbf{B}_c$  serve as the principal components for essential feature transformation, we freeze them to preserve the primary functions of the pre-trained weights. In contrast, the remaining ranks, being of lesser importance, are considered as redundant ranks  $\mathbf{A}_u$  and  $\mathbf{B}_u$ . We fine-tune them to learn the knowledge brought by the new modality. To ensure that these redundant ranks can discover a more diverse set of expertise, we divide them into  $n$  groups for subsequent orthogonal constraints. Specifically, we continue to use  $\mathbf{B}$  as the reference and perform clustering on the redundant ranks to obtain the group indices:

$$\{\mathbf{G}_1, \mathbf{G}_2, \dots, \mathbf{G}_n\} = \Gamma(\mathbf{B}_u, n), \quad (7)$$

where  $\mathbf{G}_i$  denotes the index set of the  $i$ -th group, with  $|\mathbf{G}_i| = (r - k)/n$ .  $\Gamma(\cdot)$  represents the clustering operation. We employ constrained k-means (Bradley, Bennett, and Demiriz 2000) to balance the rank capacity of each group. Then, we use  $\mathbf{G}_i$  to group the redundant ranks:

$$\begin{aligned} \mathbf{A}_{u_i} &= \{\mathbf{a}_j \mid j \in \mathbf{G}_i\}, \quad i = 1, 2, \dots, n, \\ \mathbf{B}_{u_i} &= \{\mathbf{b}_j \mid j \in \mathbf{G}_i\}, \quad i = 1, 2, \dots, n. \end{aligned} \quad (8)$$

After the completion of the redundant rank partitioning, the subsequent orthogonal constraints enforce that each group occupies non-overlapping feature spaces to achieve inter-group complementarity.

**Inter-Group Orthogonal Constraint.** Our inter-group orthogonal constraints aim to enhance the independence of feature learning among different rank groups, thereby reducing the parameter redundancy in rank spaces. Specifically,  $\mathbf{A}$  functions as a universal feature extractor, capturing the fundamental standard features within bimodal data (Tian et al. 2024). To enhance both the diversity and discriminative power of these general features, a channel orthogonality constraint is imposed on  $\mathbf{A}$ . In contrast,  $\mathbf{B}$  focuses on

learning task-specific knowledge, so rank orthogonal constraints are applied to ensure that the task knowledge carried by different ranks is complementary. The orthogonal constraint loss can be expressed as:

$$\mathcal{L}_{orth} = \sum_{i \neq j} (|\mathbf{A}_{u_i}^\top \mathbf{A}_{u_j}| + |\mathbf{B}_{u_i}^\top \mathbf{B}_{u_j}|). \quad (9)$$

Notably, calculating the above loss for all pairs of rank groups incurs a substantial computational burden. Therefore, we randomly select only one pair of different rank groups to compute the orthogonality loss at each iteration, thereby reducing the computational load.

## Objective Loss

For the tracking task, we use binary cross-entropy loss for classification and GIoU loss (Rezatofighi et al. 2019) for regression. The overall loss function is summarized as follows:

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{reg} + \lambda \cdot \mathcal{L}_{orth}, \quad (10)$$

where  $\mathcal{L}_{cls}$  represents the binary cross-entropy loss, while  $\mathcal{L}_{reg}$  represents the GIoU loss.  $\lambda$  is the balance factor.

## Experiments

### Implementation Details

We utilize LoRAT (Lin et al. 2024) as a baseline and initialize the model with its pretrained weights. As shown in Table 2, we provide two different scale variants, GOLA-B and GOLA-L. In all variants, the model trains for 10 epochs on four NVIDIA A40 GPUs, with a batch size of 128, processing 131,072 image pairs per epoch. The rank  $r$  is set to 64, the number of crucial ranks  $k$  is set to 16, and the number of redundant rank groups  $n$  is set to 8. The balance factor  $\lambda$  is set to  $1.4 \times 10^{-3}$ . The threshold  $\tau$  is set to 0.84. During the inference phase, we test the inference speed using a single NVIDIA RTX 3090 GPU.

### Comparison with State-of-the-Arts

**Evaluation on GTOT (Li et al. 2016).** GTOT includes 50 precisely paired video sequences, each with an infrared and a grayscale video. As shown in Table 1, GOLA-L establishes a new state-of-the-art performance with 95.3% MPR and 80.9% MSR. This result surpasses all previous methods and demonstrates a significant improvement over the strong baseline of CKD. GOLA-B also yields robust results, validating that our architecture exhibits strong generalizability even under the basic configuration.

**Evaluation on RGBT210 (Li et al. 2017).** RGBT210 is a precisely annotated dataset of 210 visible and infrared video pairs, comprising over 210,000 frames. As shown in Table 1, GOLA demonstrates significant advantages on RGBT210. GOLA-B achieves a PR of 90.9% and a SR of 67.0%, showing notable improvements compared to methods such as CKD and CFBT. Relative to GOLA-B, GOLA-L improves the PR and SR metrics to 92.0% and 68.7%, respectively, resulting in a further performance gain.

**Evaluation on RGBT234 (Li et al. 2019).** RGBT234 contains 234 precisely aligned visible-infrared video sequence

Methods	GTOT		RGBT210		RGBT234		LasHeR			Speed (fps)
	MPR	MSR	PR	SR	MPR	MSR	PR	NPR	SR	
APFNet (Xiao et al. 2022)	90.5	73.7	79.9	54.9	82.7	57.9	50.0	43.9	36.2	1
CAT++ (Liu et al. 2024a)	91.5	73.3	82.2	56.1	84.0	59.2	50.9	44.4	35.6	14
ViPT (Zhu et al. 2023)	-	-	-	-	83.5	61.7	65.1	-	52.5	-
TBSI (Hui et al. 2023)	-	-	85.3	62.5	87.1	63.7	69.2	65.7	55.6	36
TATrack (Wang et al. 2024)	-	-	85.3	61.8	87.2	64.4	70.2	66.7	56.1	26
STMT (Sun et al. 2024)	-	-	83.0	59.5	86.5	63.8	67.4	63.4	53.7	39
BAT (Cao et al. 2024)	-	-	-	-	86.8	64.1	70.2	-	56.3	-
CKD (Lu et al. 2024)	<b>93.2</b>	<u>77.2</u>	<u>88.4</u>	<u>65.2</u>	90.0	67.4	73.2	69.3	58.1	<u>96</u>
SDSTrack (Hou et al. 2024)	-	-	-	-	84.8	62.5	66.5	-	53.1	21
CFBT (Zeng et al. 2024)	-	-	87.7	63.0	89.9	65.9	73.2	<u>69.5</u>	58.4	-
DMD (Hu et al. 2025b)	92.4	76.8	87.0	63.7	89.3	66.7	72.6	<u>68.6</u>	57.6	17
STTrack (Hu et al. 2025a)	-	-	-	-	89.8	66.7	<u>76.0</u>	-	60.3	36
SeqTrackv2-B384 (Chen et al. 2023)	-	-	-	-	90.0	66.3	<u>71.5</u>	-	56.2	15
SUTrack-B384 (Chen et al. 2025)	-	-	-	-	<u>92.1</u>	<u>69.2</u>	75.8	-	<u>60.9</u>	32
<b>GOLA-B</b>	<u>92.8</u>	<b>78.5</b>	<b>90.9</b>	<b>67.0</b>	<b>92.2</b>	<b>69.5</b>	<b>77.5</b>	<b>73.9</b>	<b>61.6</b>	<b>125</b>
SeqTrackv2-L384 (Chen et al. 2023)	-	-	-	-	91.3	68.0	76.7	-	61.0	5
SUTrack-L384 (Chen et al. 2025)	-	-	-	-	<b>93.7</b>	<u>70.3</u>	76.9	-	<b>61.9</b>	<u>12</u>
<b>GOLA-L</b>	<b>95.3</b>	<b>80.9</b>	<b>92.0</b>	<b>68.7</b>	<u>92.8</u>	<b>71.3</b>	<b>78.1</b>	<b>74.5</b>	<b>61.9</b>	<b>64</b>

Table 1: Comparison of state-of-the-art performance across the GTOT, RGBT210, RGBT234 and LasHeR datasets. **Bold** indicates the highest performance, and underlined indicates the second-highest performance.

Variant	Template Resolution	Search Resolution	Params (M)	Trainable (%)	FLOPs (G)
GOLA-B	112×112	224×224	99	10	85
GOLA-L	112×112	224×224	336	8	284

Table 2: Details of input resolution, parameter count, and efficiency for different variants.

pairs, with roughly 234,000 frames in total. As shown in Table 1, GOLA demonstrates a significant performance advantage on RGBT234. GOLA-B performs comparably to SUTrack-L384. Furthermore, GOLA-L enhances the MPR to 92.8% and the MSR to 71.3%, with the MSR showing an improvement of 1.0% compared to SUTrack-L384. The experimental results validate the effectiveness of our approach.

**Evaluation on LasHeR (Li et al. 2021).** LasHeR comprises 1,224 visible-infrared video pairs with over 734,800 frame pairs in total. As shown in Table 1, GOLA demonstrates remarkable performance, surpassing competitive approaches such as SeqTrackv2 and SUTrack. GOLA-B achieves 77.5% of PR, 73.9% of NPR, and 61.6% of SR, outperforming larger trackers such as SeqTrackv2-L384 and SUTrack-L384. GOLA-L further elevates the PR to 78.1%, showing improvements of 1.4% and 1.2% compared to SeqTrackv2-L384 and SUTrack-L384, respectively.

**Attribute-based Performance.** LasHeR encompasses 19 distinct attributes that serve to evaluate tracker performance across various challenges as depicted in Fig. 3. GOLA consistently achieves optimal results across nearly all tested attributes. Notably, within the “HI” attribute, GOLA’s PR exceeds that of TBSI by 6.9%. Similarly, GOLA also demonstrates outstanding performance on the “HO” and “SV” at-

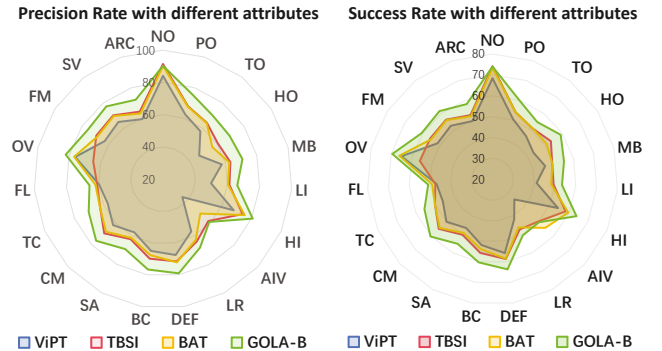


Figure 3: Comparison between GOLA-B with different trackers across various attributes in the LasHeR testing set.

tributes. These results convincingly demonstrate that GOLA is capable of acquiring diverse knowledge for various attributes, thereby enhancing tracking performance.

### Comparison between Different Fine-Tuning Methods.

Table 3 compares the performance of different fine-tuning methods. The full fine-tuning method achieves 72.5% of PR and 57.9% of SR. Fine-tuning methods such as Adapter (Houlsby et al. 2019), VPT (Jia et al. 2022), and (IA)<sup>3</sup> (Liu et al. 2022) reduce the proportion of trainable parameters, but their performance is lower compared to full fine-tuning. LoRA (Hu et al. 2022) maintains the same inference speed as the full fine-tuning method. DoRA (Liu et al. 2024b) and AdaLoRA (Zhang et al. 2023) require a longer time to fit due to the introduction of new parameters. Compared to LoRA, our method achieves significant performance improvements while reducing the number of trainable parameters by 23%.

Method	Params (M)	Trainable (%)	PR (%)	SR (%)	FLOPs (G)	Speed (fps)
FFT	88	100	72.5	57.9	77	125
Adapter	89	4	68.8	54.5	78	78
VPT	88	3	70.8	56.3	78	85
(IA) <sup>3</sup>	88	3	70.1	55.4	77	80
LoRA	99	13	76.3	60.7	85	125
AdaLoRA	99	13	57.9	44.4	85	125
DoRA	99	13	63.7	49.3	83	125
GOLA-B	99	10	77.5	61.6	85	125

Table 3: Comparison between different fine-tuning methods.

Reference	PR (%)	SR (%)
<b>A</b>	77.0	61.1
<b>B</b>	77.5	61.6

Table 4: Impact of partitioning reference on performance.

Sorting	Clustering	PR (%)	SR (%)
✓	-	77.0	61.4
-	✓	76.6	61.0
✓	✓	77.5	61.6

Table 5: Ablation on rank sorting and clustering strategy.

Type	PR (%)	SR (%)
①	76.3	60.7
②	76.8	61.2
③	76.7	61.2
GOLA-B	77.5	61.6

Table 6: Ablation of orthogonal constraint.

## Ablation Studies

**Impact of Partitioning Reference on Performance.** Table 4 demonstrates the impact of different matrices used as sorting criteria on final performance. The experimental results confirm that matrix **B** exhibits a stronger correlation with the specific task, thus enabling a more accurate identification and preservation of crucial feature information.

**Ablation on Rank Sorting and Grouping Strategy.** In Table 5, we conduct ablations on rank sorting and grouping strategies. The results indicate that rank ordering helps the model retain its generalization ability. At the same time, the grouping strategy facilitates the learning of specific knowledge within groups, thereby enhancing the expressive capability of each group. The combination of the two strategies enhances model performance more effectively.

**Effectiveness of Orthogonal Constraint.** In Table 6, we explore the effectiveness of our orthogonal constraint applied separately to **A** and **B**. ① indicates the absence of orthogonal loss, ② signifies the application of orthogonal loss solely

Pair(s)	PR (%)	SR (%)	Training Time (m)
1	77.5	61.6	110
2	77.2	61.4	120
4	77.2	61.5	150
8	77.0	61.3	200

Table 7: Ablation of sampling quantity on orthogonal loss.

Method	Params (M)	PR (%)	SR (%)	FLOPs (G)	Speed (fps)
TBSI	193	61.1	46.6	97	35
BAT	94	75.1	59.3	74	40
SeqTrackv2-B256	99	32.7	25.7	40	46
SUTrack-B224	89	71.8	56.8	70	35
GOLA-B	99	77.5	61.6	85	125

Table 8: Comparison between different trackers on DINOv2.




Figure 4: Impact of the number of crucial ranks and groups.

to the matrix **A**, and ③ denotes the application of orthogonality loss solely to the matrix **B**. The results indicate that applying the orthogonal loss simultaneously to both **A** and **B** leads to a more significant enhancement. This suggests the presence of a complementary effect between **A** and **B**, thereby validating the effectiveness of our orthogonal loss.

**Impact of Sampling Quantity.** In Table 7, we investigate the impact of the number of parameter pairs sampled randomly for the orthogonal constraint. Experimental results indicate that excessive sampling does not yield significant performance improvements. Instead, it increases computational burden. This occurs because a single set of random samples is sufficient to achieve orthogonal constraints. As feature decoupling reaches a certain extent, the effect of adding further constraints shows diminishing returns.

**Performance of Trackers using DINOv2.** To eliminate the impact of different backbones, as shown in Table 8, we apply the LoRAT pretrained DINOv2 (Oquab et al. 2024) to other methods to achieve a fairer comparison. All methods adhere to the same training protocol as our model. The new parameters introduced by TBSI and SeqTrackv2 fail to converge effectively within 10 epochs, resulting in suboptimal tracking performance. The full fine-tuning strategy employed by SUTrack leads to overfitting. Although BAT demonstrates superior tracking performance under DINOv2, its dual-stream ar-

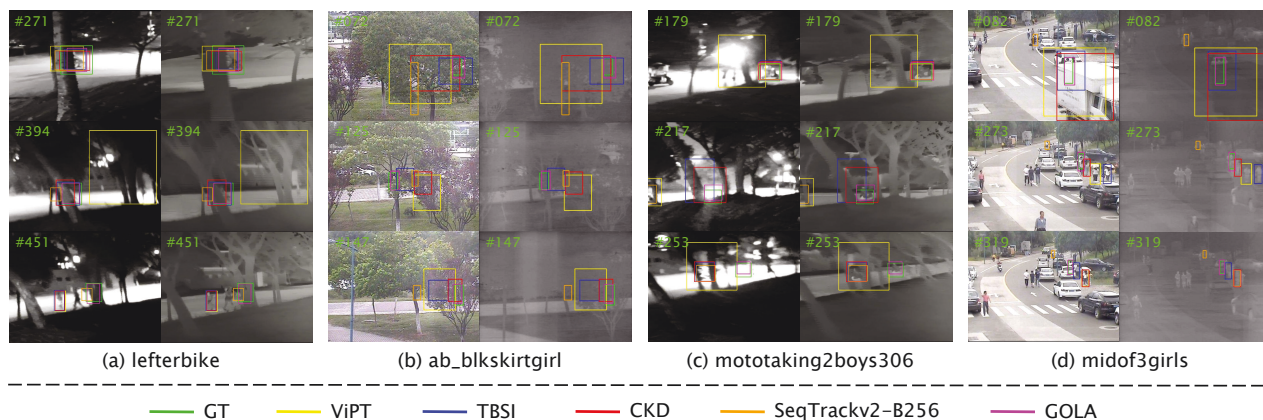


Figure 5: Qualitative comparison of GOLAB against 4 state-of-the-art trackers on 4 video sequences.




Figure 6: Visualization of t-SNE maps between rank groups. Ranks within different groups use different colors.

chitecture and additional adapter parameters limit inference speed. In contrast, GOLAB maintains optimal tracking performance while exhibiting remarkably high inference speed.

**Ablation on Crucial Rank Number.** We conduct an experiment on the number of crucial ranks  $k$  as shown in Fig. 4 (a). When  $k$  is excessively large, abundant redundant ranks become frozen, resulting in insufficient capacity for the model to learn new features. Conversely, an excessively small  $k$  fails to retain pre-trained knowledge and impairs the model’s original generalization ability.

**Ablation on Group Number.** As shown in Fig. 4 (b), we present an experiment on the number of groups  $n$ . The variable  $n$  directly influences the expressive capability of the model. As  $n$  increases, the number of ranks within each group decreases, resulting in a diminished expressive ability for individual groups. Conversely, as  $n$  decreases, having too few groups restricts the model’s ability to learn diverse knowledge, ultimately leading to a decline in performance.

## Visualization and Analysis

**Qualitative Comparison.** We qualitatively compare our method with other RGB-T trackers. As shown in Fig. 5, we select four representative sequences from LasHeR, covering diverse tracking challenges such as partial and total occlusions, low illumination, and similar-object interference,

to assess performance in realistic scenarios visually. Taking the third sequence as an example, our method demonstrates robust tracking capabilities even in the presence of severe dynamic blur and varying illumination. Results indicate that GOLAB learns specific knowledge tailored to different challenges, thereby enhancing tracking performance.

**Visualization of Rank Groups.** We visualize the ranks of different groups using t-SNE in Fig. 6. For LoRA, we identify the corresponding ranks for visualization by utilizing the grouping indices from GOLAB. Compared to LoRA, the rank feature points within the same group of GOLAB exhibit a high degree of clustering. This suggests intra-group ranks form a stronger synergistic effect during learning, enabling concentrated expression around specific tasks. Rank orders between different groups show distinct separation, indicating that the groups have learned complementary expressive capabilities. Visualization results further validate the superiority of our approach in learning strategies.

## Conclusion

In this paper, we propose a Group Orthogonal Low-Rank Adaptation (GOLAB) framework for RGB-T tracking, aiming to address the limitations in parameter redundancy and low-rank space expression capabilities. GOLAB identifies and retains crucial ranks by introducing a rank decomposition partitioning strategy, thereby preserving the priors of the pre-trained weights while clustering redundant ranks for structured learning. Furthermore, we implement an inter-group orthogonal constraint strategy to mitigate information redundancy within low-rank spaces, enhancing the model’s expressive capabilities for addressing diverse challenges. Compared to full fine-tuning, GOLAB reduces the number of trainable parameters to alleviate overfitting and achieves parameter merging during inference without incurring additional delays. Experimental results validate the effectiveness and efficiency of GOLAB, demonstrating that our method significantly outperforms state-of-the-art methods while maintaining real-time performance. We hope that GOLAB can provide a new perspective for fine-tuning methods in RGB-T tracking and facilitate further advancements in this field.

## Acknowledgments

This work was supported by the Joint Funds of the National Natural Science Foundation of China (U2441251, U24A20218), in part by the National Natural Science Foundation of China (62303046).

## References

- Bradley, P.; Bennett, K.; and Demiriz, A. 2000. Constrained K-Means Clustering.
- Cao, B.; Guo, J.; Zhu, P.; and Hu, Q. 2024. Bi-directional adapter for multimodal tracking. In *AAAI*, volume 38, 927–935.
- Chen, X.; Kang, B.; Geng, W.; Zhu, J.; Liu, Y.; Wang, D.; and Lu, H. 2025. Sutrack: Towards simple and unified single object tracking. In *AAAI*, volume 39, 2239–2247.
- Chen, X.; Kang, B.; Zhu, J.; Wang, D.; Peng, H.; and Lu, H. 2023. Unified Sequence-to-Sequence Learning for Single- and Multi-Modal Visual Object Tracking. *arXiv preprint arXiv:2304.14394*.
- Gao, C.; Chen, K.; Rao, J.; Sun, B.; Liu, R.; Peng, D.; Zhang, Y.; Guo, X.; Yang, J.; and Subrahmanian, V. 2024. Higher layers need more lora experts. *arXiv preprint arXiv:2402.08562*.
- Hou, X.; Xing, J.; Qian, Y.; Guo, Y.; Xin, S.; Chen, J.; Tang, K.; Wang, M.; Jiang, Z.; Liu, L.; et al. 2024. Sdstrack: Self-distillation symmetric adapter learning for multi-modal visual object tracking. In *CVPR*, 26551–26561.
- Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; and Gelly, S. 2019. Parameter-efficient transfer learning for NLP. In *ICML*, 2790–2799. PMLR.
- Hu, E. J.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*.
- Hu, X.; Tai, Y.; Zhao, X.; Zhao, C.; Zhang, Z.; Li, J.; Zhong, B.; and Yang, J. 2025a. Exploiting multimodal spatial-temporal patterns for video object tracking. In *AAAI*, volume 39, 3581–3589.
- Hu, Y.; Shao, Z.; Fan, B.; and Liu, H. 2025b. Dual-Level Modality De-Biasing for RGB-T Tracking. *TIP*.
- Hui, T.; Xun, Z.; Peng, F.; Huang, J.; Wei, X.; Wei, X.; Dai, J.; Han, J.; and Liu, S. 2023. Bridging search region interaction with template for rgb-t tracking. In *CVPR*, 13630–13639.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *ECCV*, 709–727. Springer.
- Jie, S.; Deng, Z.-H.; Chen, S.; and Jin, Z. 2024. Convolutional bypasses are better vision transformer adapters. In *ECAI*, 202–209. IOS Press.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In *EMNLP*, 3045–3059.
- Li, C.; Cheng, H.; Hu, S.; Liu, X.; Tang, J.; and Lin, L. 2016. Learning collaborative sparse representation for grayscale-thermal tracking. *TIP*, 25(12): 5743–5756.
- Li, C.; Liang, X.; Lu, Y.; Zhao, N.; and Tang, J. 2019. RGB-T object tracking: Benchmark and baseline. *PR*, 96: 106977.
- Li, C.; Xue, W.; Jia, Y.; Qu, Z.; Luo, B.; Tang, J.; and Sun, D. 2021. LasHeR: A large-scale high-diversity benchmark for RGBT tracking. *TIP*, 31: 392–404.
- Li, C.; Zhao, N.; Lu, Y.; Zhu, C.; and Tang, J. 2017. Weighted sparse representation regularized graph learning for RGB-T object tracking. In *ACM MM*, 1856–1864.
- Li, D.; Ma, Y.; Wang, N.; Cheng, Z.; Duan, L.; Zuo, J.; Yang, C.; and Tang, M. 2024. MixLoRA: Enhancing Large Language Models Fine-Tuning with LoRA based Mixture of Experts. *CoRR*.
- Li, X. L.; and Liang, P. 2021. Prefix-Tuning: Optimizing Continuous Prompts for Generation. In *ACL*, 4582–4597.
- Lin, L.; Fan, H.; Zhang, Z.; Wang, Y.; Xu, Y.; and Ling, H. 2024. Tracking meets lora: Faster training, larger model, stronger performance. In *ECCV*, 300–318. Springer.
- Liu, H.; Tam, D.; Muqeeth, M.; Mohta, J.; Huang, T.; Bansal, M.; and Raffel, C. A. 2022. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning. *NeurIPS*, 35: 1950–1965.
- Liu, L.; Li, C.; Xiao, Y.; Ruan, R.; and Fan, M. 2024a. RGBT Tracking via Challenge-Based Appearance Disentanglement and Interaction. *TIP*, 33: 1753–1767.
- Liu, S.-Y.; Wang, C.-Y.; Yin, H.; Molchanov, P.; Wang, Y.-C. F.; Cheng, K.-T.; and Chen, M.-H. 2024b. Dora: Weight-decomposed low-rank adaptation. In *ICML*.
- Lu, A.; Zhao, J.; Li, C.; Xiao, Y.; and Luo, B. 2024. Breaking modality gap in RGBT tracking: Coupled knowledge distillation. In *ACM MM*, 9291–9300.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H. V.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; HAZIZA, D.; Massa, F.; El-Nouby, A.; et al. 2024. DINOv2: Learning Robust Visual Features without Supervision. *TMLR*.
- Rezatofighi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I.; and Savarese, S. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *CVPR*, 658–666.
- Shao, Z.; Hu, Y.; Fan, B.; and Liu, H. 2025. PURA: Parameter Update-Recovery Test-Time Adaption for RGB-T Tracking. In *CVPR*, 22089–22098.
- Sun, D.; Pan, Y.; Lu, A.; Li, C.; and Luo, B. 2024. Transformer RGBT Tracking With Spatio-Temporal Multimodal Tokens. *TCSVT*, 34(11\_Part\_2): 12059–12072.
- Tian, C.; Shi, Z.; Guo, Z.; Li, L.; and Xu, C.-Z. 2024. Hydralora: An asymmetric lora architecture for efficient fine-tuning. *NeurIPS*, 37: 9565–9584.
- Wang, H.; Liu, X.; Li, Y.; Sun, M.; Yuan, D.; and Liu, J. 2024. Temporal adaptive rgbt tracking with modality prompt. In *AAAI*, volume 38, 5436–5444.
- Wu, Z.; Zheng, J.; Ren, X.; Vasluianu, F.-A.; Ma, C.; Paudel, D. P.; Van Gool, L.; and Timofte, R. 2024. Single-model and any-modality for video object tracking. In *CVPR*, 19156–19166.

Xia, J.; Shi, D.; Song, K.; Song, L.; Wang, X.; Jin, S.; Zhao, C.; Cheng, Y.; Jin, L.; Zhu, Z.; et al. 2024. Unified Single-Stage Transformer Network for Efficient RGB-T Tracking. In *IJCAI*.

Xiao, Y.; Yang, M.; Li, C.; Liu, L.; and Tang, J. 2022. Attribute-based progressive fusion network for rgbt tracking. In *AAAI*, volume 36, 2831–2838.

Zeng, Z.; Liu, X.; Sun, M.; Wang, H.; and Liu, J. 2024. Cross Fusion RGB-T Tracking with Bi-directional Adapter. *CoRR*.

Zhang, Q.; Chen, M.; Bukharin, A.; He, P.; Cheng, Y.; Chen, W.; and Zhao, T. 2023. Adaptive Budget Allocation for Parameter-Efficient Fine-Tuning. In *ICLR*. Openreview.

Zhu, J.; Lai, S.; Chen, X.; Wang, D.; and Lu, H. 2023. Visual prompt multi-modal tracking. In *CVPR*, 9516–9526.