

# General Compression Framework for Efficient Transformer Object Tracking

Lingyi Hong<sup>1</sup> Jinglun Li<sup>2</sup> Xinyu Zhou<sup>1</sup> Shilin Yan<sup>1</sup> Pinxue Guo<sup>2</sup> Kaixun Jiang<sup>2</sup> Zhaoyu Chen<sup>2</sup>  
Shuyong Gao<sup>1</sup> Runze Li<sup>3</sup> Xingdong Sheng<sup>3</sup> Wei Zhang<sup>1\*</sup> Hong Lu<sup>1\*</sup> Wenqiang Zhang<sup>1,2\*</sup>

<sup>1</sup> Shanghai Key Lab of Intelligent Information Processing,

College of Computer Science and Artificial Intelligence, Fudan University

<sup>2</sup> College of Intelligent Robotics and Advanced Manufacturing, Fudan University

<sup>3</sup> Lenovo Research

honglyhly@gmail.com, {weizh, honglu, wqzhang}@fudan.edu.cn

## Abstract

*Previous works have attempted to improve tracking efficiency through lightweight architecture design or knowledge distillation from teacher models to compact student trackers. However, these solutions often sacrifice accuracy for speed to a great extent, and also have the problems of complex training process and structural limitations. Thus, we propose a general model compression framework for efficient transformer object tracking, named CompressTracker, to reduce model size while preserving tracking accuracy. Our approach features a novel stage division strategy that segments the transformer layers of the teacher model into distinct stages to break the limitation of model structure. Additionally, we also design a unique replacement training technique that randomly substitutes specific stages in the student model with those from the teacher model, as opposed to training the student model in isolation. Replacement training enhances the student model's ability to replicate the teacher model's behavior and simplifies the training process. To further forcing student model to emulate teacher model, we incorporate prediction guidance and stage-wise feature mimicking to provide additional supervision during the teacher model's compression process. CompressTracker is structurally agnostic, making it compatible with any transformer architecture. We conduct a series of experiment to verify the effectiveness and generalizability of our CompressTracker. Our CompressTracker-SUTrack, compressed from SUTrack, retains about 99% performance on LaSOT (72.2% AUC) while achieves 2.42× speed up. Code is available at [here](#).*

## 1. Introduction

Visual object tracking is tasked with continuously localizing a target object across video frames based on the initial

bounding box in the first frame. Transformer-based trackers have achieved promising performance on well-established benchmarks, their deployment on resource-restricted device remains a significant challenge. Developing a strong tracker with high efficiency is of great significance.

Previous works attempt to reduce model inference cost through lightweight tracker design or knowledge transfer from teacher to student models. However, these approaches face several limitations. (1) **Inferior Accuracy.** Certain works propose lightweight tracking models [4, 6, 11, 21, 31] or employ neural architecture search (NAS) [47] often suffer from underfitting and inferior performance because of the limited number of parameters. (2) **Complex Training.** Some works [15] transfers the knowledge from a teacher tracker to a student model, while effective, requiring complex multi-stage and time-consuming training procedures and resulting in suboptimal performance. (3) **Structure Limitation.** The model reduction paradigm in [15] severely restricts the structure of student to be consistent with teacher.

Thus, we introduce CompressTracker, a novel and general model compression framework to enhance the efficiency of transformer tracking models, which is built upon three interconnected innovations. First, to address the structure limitation, we introduce the **stage division** strategy. Current dominant trackers are one-stream models [4, 11, 15, 49] with sequential transformer encoder layers that refine temporal matching features. Given this layer-wise refinement, it becomes a natural progression to consider the model not as a single entity but as a series of interconnected stages and encourage student tracker to align teacher model at each stage. Stage division partitions the teacher model into distinct stages that correspond to the layers of a simpler student model. This is achieved by dividing the teacher model into a number of stages equivalent to the student model's layers. Each stage in the student model learns to replicate the functional behavior of its corresponding teacher stage. This division is not merely a structural alteration but a strategic

\*Corresponding Author

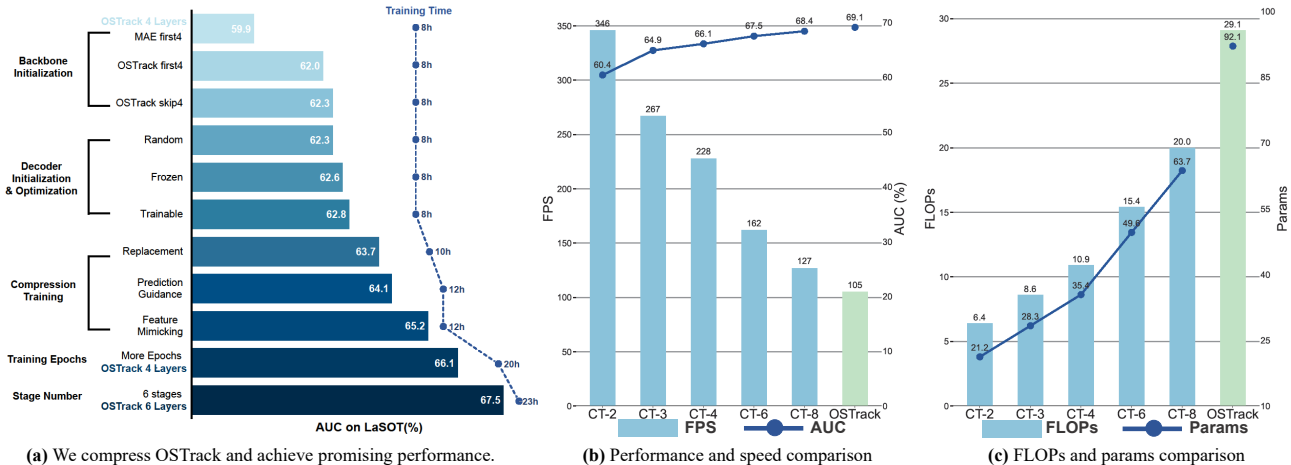


Figure 1. We apply our framework to OSTRack under several different layer configurations. (a) We implement each enhancement into our CompressTracker step by step. The training time is calculated by using 8 NVIDIA RTX 3090 GPUs. Notably, our CompressTracker-4 accelerates OSTRack by  $2.17\times$  while preserving approximately  $96\%$  of its original accuracy, thereby demonstrating the effectiveness of our framework. (b) Performance and speed comparison of CompressTracker variants with different numbers of layers. CT-x refers to a version of CompressTracker with 'x' layers. (c) FLOPs and parameters comparison of CompressTracker variants with different numbers of layers.

educational approach, which allows student to grasp not only what to track but also how, facilitating a more effective and fine-grained compression process.

Secondly, the **stage division** lays the foundation for our **replacement training**, which tackles the complex training challenge by dynamically integrating teacher and student models during training. Replacement training strategically intertwines the teacher and student models. The core of this methodology is the dynamic substitution of stages during training. we randomly select stages from the student model and replace them with the corresponding stages from the teacher model. Replacement training permits the unaltered stages of the teacher model to collaboratively inform and enhance the learning of the substituted stages in the student model rather than supervising the entire student model as a single entity. The student model is not merely learning in parallel but is directly engaging with the teacher's learned behaviors. After training, we can just combine each stage of student model for inference. The replacement training leads to a more authentic replication of the teacher's tracking strategies and helps to prevent the student model from overfitting to specific stages of the teacher model.

Thirdly, to further improve accuracy, we introduce **prediction guidance** and **stage-wise feature mimicking**. By using the predictions of the teacher model as a reference, the student model can converge more quickly. Feature mimicking systematically aligns the feature representations learned at each stage of the student model with those of the teacher model, thereby promoting a more accurate and consistent learning. Each of these three innovations is sequentially built upon the previous one, forming a cohesive and highly effective compression framework that maximizes efficiency while preserving performance. In Figure 1 (a), we show procedure

and results we are able to achieve with each step toward an efficient transformer tracker.

Compared to previous works, our CompressTracker holds many merits. (1) **Enhanced Mimicking and Performance.** CompressTracker enables the student model to better mimic the teacher model, resulting in better performance. As shown in Figure 1, our CompressTracker achieves an optimal balance between accuracy and inference speed. (2) **Simplified Training Process.** Our CompressTracker streamlines training into a single but efficient step. This simplification not only reduces the time and resources required for training but also minimizes the potential for sub-optimal performance associated with complex procedures. The training process for CompressTracker-4 requires merely 20 hours on 8 NVIDIA RTX 3090 GPUs. (3) **Heterogeneous Model Compression.** Our stage division strategy gives a high degree of flexibility in the design of the student model. Our framework supports any transformer architecture for student model, which is not restricted to the same structure of teacher tracker. The number of layers and their structure are not predetermined but can be tailored to fit the specific computational constraints and requirements of the deployment environment.

Our contribution can be summarized as follows: (1) We introduce a novel and general model compression framework, CompressTracker, to facilitate the efficient transformer-based object tracking. (2) We propose a stage division strategy that enables a fine-grained imitation of the teacher model at the stage level, enhancing the precision and efficiency of knowledge transfer. (3) We propose the replacement training to improve the student model's capacity to replicate the teacher model's behavior. (4) We further incorporate the prediction guidance and feature mimicking to accelerate and refine the learning process of the student model. (5)

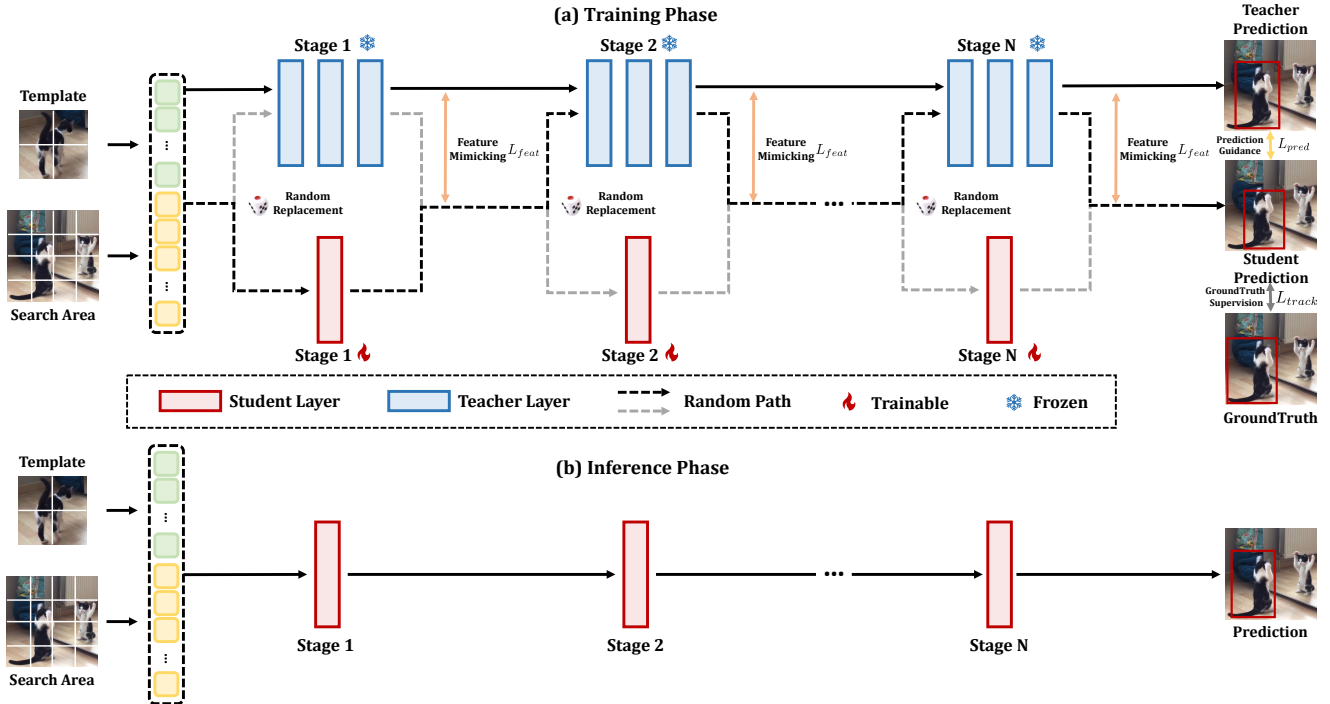


Figure 2. **CompressTracker Framework.** (a) In the training phase, we divide both the teacher model and student model into an identical number of stages. We implement a series of training strategies including replacement training, prediction guidance, and stage-wise feature mimicking, to enhance the student model’s ability to emulate the teacher model. The dotted lines represent the randomly selected paths for replacement training, with black dotted lines indicating the chosen path, while gray dotted lines denote paths not selected in a specific training iteration. (b) During inference process, we simply combine each stage of the student model for testing purposes.

Our CompressTracker breaks structural limitations, adapting to various transformer architectures for student model. Our CompressTracker-SUTrack outperforms existing models, notably accelerating SUTrack [49] by  $2.42\times$  while preserving approximately 99% accuracy (72.2% AUC on LaSOT).

## 2. Related Work

**Visual Object Tracking.** Visual object tracking aims to localize the target object of each frame based on its initial appearance. Previous tracking methods [2, 5, 10, 16, 23, 25, 27, 32, 33, 46, 51] utilize a two-stream pipeline to decouple the feature extraction and relation modeling. Recently, the one-stream pipeline hold the dominant role. [1, 8, 12, 14, 15, 19, 24, 26, 28, 42, 49, 53, 54] combine feature extraction and relation modeling into a unified process. These models are built upon vision transformer, which consists of a series of transformer encoder layers. Thanks to a more adequate relationship modeling between template and search frame, one-stream models achieve impressive performance. However, these models suffer from low inference efficiency, which is the main obstacle to practical deployment.

**Efficient Tracking.** Some works have attempted to speed up tracking models. [47] utilizes neural architecture search (NAS) to search a light Siamese network, and the search-

ing process is complex. [4, 6, 11, 31] design a lightweight tracking model, but the small number of parameters restricts the accuracy to a large degree. MixFormerV2 [15] propose a complex multi-stage model reduction strategy. Although MixFormerV2-S achieves real-time speed on CPU, the multi-stage training strategy is time consuming, which requires about 120 hours (5 days) on 8 Nvidia RTX8000 GPUs, even several times the original training time of MixFormer [14]. Any suboptimal performance during these stages impact the final model’s performance negatively. Besides, the reduction paradigm imposes constraints on the design of student models. To address these shortcuts, we propose the general model compression framework, CompressTracker, to explore the roadmap toward an end-to-end and training-efficient model compression for lightweight transformer-based tracker. Our CompressTracker break the structure restriction and achieves balance between speed and accuracy.

**Transformer Compression.** Model compression aims to reduce the size and computational cost of a large model while retaining as much performance as possible, and recently many attempts have been made to speed up a large pretrained transformer model. [18] reduced the number of parameters through pruning technique, and [40] accomplished the quantization of BERT to 2-bits utilizing Hessian information. [30, 39, 41, 43] leverage the knowledge dis-

Method	LaSOT			LaSOT <sub>ext</sub>		TNL2K		TrackingNet			UAV123		FPS	
	AUC	P <sub>Norm</sub>	P	AUC	P	AUC	P	AUC	P <sub>Norm</sub>	P	AUC	P	GPU	CPU
<b>CompressTracker-OSTrack</b> [49]	66.1 <sub>96%</sub>	75.2	70.6	45.7 <sub>96%</sub>	50.8	53.6 <sub>99%</sub>	52.5	82.1 <sub>99%</sub>	87.6	80.1	67.4 <sub>99%</sub>	88.0	<b>228</b> <sub>2.17×</sub>	18 <sub>2.11×</sub>
<b>CompressTracker-SeqTrack</b> [12]	68.1 <sub>95%</sub>	77.3	73.4	47.9 <sub>96%</sub>	56.4	54.5 <sub>99%</sub>	54.2	83.1 <sub>98%</sub>	88.1	81.7	68.4 <sub>98%</sub>	90.1	78 <sub>1.95×</sub>	10 <sub>2.46×</sub>
<b>CompressTracker-ODTrack</b> [52]	70.5 <sub>96%</sub>	79.5	76.9	50.9 <sub>97%</sub>	60.8	58.2 <sub>99%</sub>	58.9	82.8 <sub>97%</sub>	87.9	81.4	69.2 <sub>98%</sub>	90.8	87 <sub>2.71×</sub>	13 <sub>3.18×</sub>
<b>CompressTracker-SUTrack</b> [13]	<b>72.2</b> <sub>99%</sub>	<b>80.7</b>	<b>78.2</b>	<b>52.1</b> <sub>98%</sub>	<b>63.2</b>	<b>64.4</b> <sub>99%</sub>	<b>69.9</b>	<b>84.3</b> <sub>99%</sub>	<b>88.2</b>	<b>83.3</b>	<b>71.0</b> <sub>99%</sub>	<b>91.4</b>	134 <sub>2.42×</sub>	<b>24</b> <sub>2.34×</sub>
SUTrack-T [13]	69.6	79.3	75.4	50.2	57.0	-	-	82.7	87.2	80.8	69.4	-	100	23
HiT-Base [31]	64.6	73.3	68.1	44.1	-	-	-	80.0	84.4	77.3	65.6	-	175	33
HiT-Samll [31]	60.5	68.3	61.5	40.4	-	-	-	77.7	81.9	73.1	63.3	-	192	72
HiT-Tiny [31]	54.8	60.5	52.9	35.8	-	-	-	74.6	78.1	68.8	53.2	-	204	76
SMAT [21]	61.7	71.1	64.6	-	-	-	-	78.6	84.2	75.6	64.3	83.9	158	33
MixFormerV2-S [15]	60.6	69.9	60.4	43.6	46.2	48.3	43.0	75.8	81.1	70.4	65.8	86.8	325	30
FEAR-L [6]	57.9	68.6	60.9	-	-	-	-	-	-	-	-	-	-	-
FEAR-XS [6]	53.5	64.1	54.5	-	-	-	-	-	-	-	-	-	80	60
HCAE [11]	59.0	68.3	60.5	-	-	-	-	76.6	82.6	72.9	63.6	-	195	45
E.T.Track [4]	59.1	-	-	-	-	-	-	74.5	80.3	70.6	62.3	-	150	47
LightTrack-LargeA [47]	55.5	-	56.1	-	-	-	-	73.6	78.8	70.0	-	-	-	-
LightTrack-Mobile [47]	53.8	-	53.7	-	-	-	-	72.5	77.9	69.5	-	-	120	41
STARK-Lightning [46]	58.6	69.0	57.9	-	-	-	-	-	-	-	-	-	200	-
DiMP [3]	56.9	65.0	56.7	-	-	-	-	74.0	80.1	68.7	65.4	-	77	-
SiamFC++ [44]	54.4	62.3	54.7	-	-	-	-	75.4	80.0	70.5	-	-	90	-

Table 1. **Compress Object Tracker.** We compress several trackers such as SUTrack [13], ODTrack [52], and so on, into students with an optimal trade-off between accuracy and efficiency. We report the performance on 5 benchmarks and calculate the performance gap in comparison to the origin teacher model. CompressTracker outperforms previous models in both speed and accuracy.

tillation to transfer the knowledge from teacher to student model and exploit pretrained model. Beyond language models, considerable focus has also been placed on compressing vision transformer models. [7, 9, 20, 38, 45, 48, 50] utilize multiple model compression techniques to compress vision transformer models. MixFormerV2 [15] proposed a two-stage model reduction paradigm to distill a lightweight tracker, relying on the complex multi-stage distillation training. However, our CompressTracker propose an end-to-end and efficient compression training to achieve any transformer structure compression, which speed up SUTrack **2.42**× while maintaining about **99%** accuracy.

### 3. CompressTracker

In this section, we will introduce our proposed general model compression framework, CompressTracker. The workflow of our CompressTracker in illustrated in Figure 2.

#### 3.1. Stage Division

Recently, transformer-based one-stream tracking models [8, 14, 15, 49] have dominated visual object tracking, consisting of multiple transformer encoder layers that progressively refine temporal matching features. Building upon this layer-wise refinement, we introduce the stage division strategy, which segments the model into a series of sequential stages. This approach encourages the student model to emulate the teacher model’s behavior at each individual stage. Specifically, we denote the pretrained tracker and the compressed model as *teacher* and *student* model, with  $N_t$  and  $N$  layers, respectively. Both teacher and student models are then divided into  $N$  stages, where each student stage encompasses a single layer, and each corresponding stage in teacher may

aggregate multiple layers, which can be formulated as:

$$teacher = \{stage_1^t, stage_2^t, \dots, stage_N^t\}, \quad (1)$$

$$student = \{stage_1^s, stage_2^s, \dots, stage_N^s\}, \quad (2)$$

where  $stage_i^t$  and  $stage_i^s$  denote the corresponding stage  $i$  in teacher and student model, respectively. For a specific stage  $i$ , we establish a correspondence between the stages of the teacher and student models. The objective of stage division is to enforce each stage of the student model to replicate its counterpart in the teacher model. This stage division strategy breaks the traditional approach that treats the model as an indivisible whole [4, 6, 11, 15]. Instead, it enables a fine-grained learning process where the student model transfers knowledge from the teacher in a more detailed, stage-specific manner.

Unlike the reduction paradigm adopted in [15], our CompressTracker framework facilitates support for arbitrary transformer structures of the student tracker, thanks to our innovative stage-wise division design. To align the size and channel dimensions of the student model’s temporal matching features with those of the teacher model, we implement input and output projection layers before and after the student layers, respectively. These layers ensure compatibility between teacher and student models, allowing greater flexibility in the student model’s architecture. During inference, these projection layers are omitted.

#### 3.2. Replacement Training

During training, we adopt the replacement training to integrates teacher model and student models, diverging from the conventional practice of training the student model in isolation. In each iteration, a stochastic process determines which student model stages are replaced by corresponding

Method	LaSOT			LaSOT <sub>ext</sub>		TNL2K		TrackingNet			UAV123		FPS
	AUC	P <sub>Norm</sub>	P	AUC	P	AUC	P	AUC	P <sub>Norm</sub>	P	AUC	P	
<i>Model Generalization</i>													
<b>CompressTracker-OSTrack [49]</b>	66.1 <sub>96%</sub>	75.2	70.6	45.7 <sub>96%</sub>	50.8	53.6 <sub>99%</sub>	52.5	82.1 <sub>99%</sub>	87.6	80.1	67.4 <sub>99%</sub>	88.0	228 <sub>2.17×</sub>
<b>CompressTracker-SeqTrack [12]</b>	68.1 <sub>95%</sub>	77.3	73.4	47.9 <sub>96%</sub>	56.4	54.5 <sub>99%</sub>	54.2	83.1 <sub>98%</sub>	88.1	81.7	68.4 <sub>98%</sub>	90.1	78 <sub>1.95×</sub>
<b>CompressTracker-ODTrack [52]</b>	70.5 <sub>96%</sub>	79.5	76.9	50.9 <sub>97%</sub>	60.8	58.2 <sub>99%</sub>	58.9	82.8 <sub>97%</sub>	87.9	81.4	69.2 <sub>98%</sub>	90.8	87 <sub>2.71×</sub>
<b>CompressTracker-SUTrack [13]</b>	72.2 <sub>99%</sub>	80.7	78.2	52.1 <sub>98%</sub>	63.2	64.4 <sub>99%</sub>	69.9	84.3 <sub>99%</sub>	88.2	83.3	71.0 <sub>99%</sub>	91.4	134 <sub>2.42×</sub>
<i>Stage Scalability</i>													
<b>CompressTracker-OSTrack-2</b>	60.4 <sub>87%</sub>	68.5	61.5	40.4 <sub>85%</sub>	43.8	48.5 <sub>89%</sub>	45.0	78.2 <sub>94%</sub>	83.3	74.8	62.5 <sub>92%</sub>	82.5	346 <sub>3.30×</sub>
<b>CompressTracker-OSTrack-3</b>	64.9 <sub>94%</sub>	74.0	68.4	44.6 <sub>94%</sub>	49.6	52.6 <sub>97%</sub>	50.9	81.6 <sub>98%</sub>	86.7	79.4	65.4 <sub>96%</sub>	88.3	267 <sub>2.54×</sub>
<b>CompressTracker-OSTrack-4</b>	66.1 <sub>96%</sub>	75.2	70.6	45.7 <sub>96%</sub>	50.8	53.6 <sub>99%</sub>	52.5	82.1 <sub>99%</sub>	87.6	80.1	67.4 <sub>99%</sub>	88.0	228 <sub>2.17×</sub>
<b>CompressTracker-OSTrack-6</b>	67.5 <sub>98%</sub>	77.5	72.4	46.7 <sub>98%</sub>	52.5	54.7 <sub>101%</sub>	54.3	82.9 <sub>99%</sub>	87.8	81.5	67.9 <sub>99%</sub>	88.7	162 <sub>1.54×</sub>
<b>CompressTracker-OSTrack-8</b>	68.4 <sub>99%</sub>	78.0	73.1	47.2 <sub>99%</sub>	53.1	55.2 <sub>102%</sub>	54.8	83.3 <sub>101%</sub>	88.0	81.9	68.2 <sub>99%</sub>	89.0	127 <sub>1.21×</sub>
<i>Larger Transformer Scalability</i>													
<b>CompressTracker-OSTrack-L-4</b>	67.5 <sub>96%</sub>	77.9	72.7	45.9 <sub>98%</sub>	51.5	58.3 <sub>98%</sub>	58.5	83.2 <sub>99%</sub>	87.7	81.2	67.4 <sub>99%</sub>	88.2	228 <sub>2.84×</sub>
<i>Higher Resolution Scalability</i>													
<b>CompressTracker-OSTrack-384-4</b>	67.7 <sub>96%</sub>	78.3	73.6	48.1 <sub>96%</sub>	58.0	54.3 <sub>99%</sub>	53.7	82.7 <sub>99%</sub>	87.3	81.0	68.2 <sub>98%</sub>	88.8	228 <sub>3.90×</sub>

Table 2. **Generalization of CompressTracker.** Our CompressTracker can be applied to any teacher model, any resolution, any stage number, any teacher model size, and any student architectures. CompressTracker-m-x denotes the compressed student model with 'x' layers with 'm' as teacher model. We report the performance on 5 benchmarks and calculate the performance gap in comparison to the origin teacher model.

teacher stages. For the specific stage  $i$ , the forward process in conventional isolated propagation can be described as:

$$h_i = \text{stage}_i^s(h_{i-1}), \quad (3)$$

where  $h_{i-1}$  is the input of the  $i$  student stage. However, in our replacement training, we decide whether to replace or not by random Bernoulli sampling  $b_i$  with probability  $p$ , where  $b_i \in \{0, 1\}$ . If  $b_i$  equals 1, the output from the preceding  $i - 1$ -th stage is directed to the  $i$  student stage, otherwise, we channel the output into the  $i$  frozen teacher stage, which can be formulated as:

$$h_i = \begin{cases} \text{stage}_i^t(h_{i-1}), & r_i = 0, \\ \text{stage}_i^s(h_{i-1}), & r_i = 1, \end{cases} \quad r_i \sim \text{Bernoulli}(p). \quad (4)$$

This replacement training creates a collaborative learning environment where the teacher model dynamically supervises the student model. The unreplaced stages of teacher provide valuable contextual supervision for a specific stage in the student model. Consequently, the student model is not operating in parallel but is actively engaged with and learning from the teacher’s established behaviors. For the optimization of student model, we only require the groundtruth box and denote the loss as  $L_{\text{track}}$ . After training, the student model’s stages are harmoniously combined for inference.

### 3.3. Prediction Guidance & Stage-wise Feature Mimicking

Replacement training enables the student model to learn the behavior of each individual stage, resulting in enhanced performance. However, mimicking the teacher directly can be too difficult for a smaller student model. Thus, we employ the teacher’s predictions to further guide the learning of compressed tracker. We apply the same loss as  $L_{\text{track}}$  for

prediction guidance, which is denoted as  $L_{\text{pred}}$ . This guidance accelerates and stabilizes the student’s learning, helping it assimilate knowledge from the teacher more effectively.

While prediction guidance accelerates the convergence, the student tracker might not entirely match the complex behavior of the teacher model. We introduce the stage-wise feature mimicking to further synchronize the temporal matching features between corresponding stages of the teacher and student models. This alignment is quantified by calculating the  $L_2$  distance between the outputs of these stages, which is referred as  $L_{\text{feat}}$ . It is worth noting that any metric assessing the discrepancy in feature distributions can serve as the loss function. However, we choose a simple  $L_2$  distance rather than a complex loss to highlight the effectiveness of our stage division and replacement training strategies. The stage-wise feature mimicking both promotes a closer similarity in feature representations of corresponding stages and enhances overall coherence between teacher and student.

### 3.4. Progressive Replacement

In Section 3.2, we describe the replacement training strategy. Although setting the Bernoulli sampling probability  $p$  as a constant value can realize the compression, these stages have not been trained together at the same time and there may be some dissonance. A further finetuning step is necessary to achieve better harmony among the stages. Thus, we introduce a progressive replacement strategy to bridges the gap between the two initially separate training phases, fostering an end-to-end easy-to-hard learning process. By adjusting the value of  $p$ , we can control the number of stages to be replaced. The value of  $p$  gradually increases from  $p_{\text{init}}$  to 1.0, allowing for a more incremental and coherent training:

#	Init. method	AUC
1	MAE-first4	59.9%
2	OSTrack-first4	62.0%
3	<b>OSTrack-skip4</b>	<b>62.3%</b>

Table 3. **Backbone Initialization.** 'MAE-first4' denotes initializing the student model using the first 4 layers of MAE-B. 'OSTrack-skip4' represents utilizing every fourth layer of OSTrack for the student model.

#	Layer Split	AUC
1	<b>Even</b>	<b>62.8%</b>
2	Uneven	62.7%

Table 5. **Stage Division.** 'Even' denotes evenly dividing stage, and 'Uneven' means that the layer number of each stage in teacher model is 2,2,6,2.

#	Epochs	AUC
1	300	65.2%
2	<b>500</b>	<b>66.1%</b>

Table 8. **Training Epochs.** '300' and '500' denote total epochs.

#	Replacement	AUC	Training Time
1	<b>Random</b>	<b>65.2%</b>	<b>12 h</b>
2	Decouple-300	64.6%	16 h

Table 6. **Replacement training.** 'Random' denotes our replacement training, and 'Decouple-300' represents decoupling the training of each stage.

#	Model	Training Time
1	<b>CompressTracker-4</b>	<b>20 h</b>
2	OSTrack	17 h
3	MixFormerV2-S	120 h

Table 9. **Training Time** comparison with other methods.

#	Init. & Opt.	AUC
1	Random & Trainable	62.3%
2	Teacher & Frozen	62.6%
3	<b>Teacher &amp; Trainable</b>	<b>62.8%</b>

Table 4. **Decoder Initialization and Optimization.** 'Random' denotes randomly initialized decoder, and 'Teacher' means the decoder is initialized with teacher parameters. 'Frozen' represents that the decoder is frozen, and 'Trainable' denotes decoder is trainable.

#	Replacement	AUC
1	<b>w/ Progressive</b>	<b>65.2%</b>
2	w/o Progressive	64.8%

Table 7. **Progressive Replacement.** We conduct additional finetune if progressive replacement is disabled.

#	Method	AUC	FPS
1	<b>CompressTracker-4</b>	66.1	228
2	Distillation	63.8	228
3	Pruning (MixFormerV2-S)	60.6	325

Table 10. **Comparison with Other Compression Techniques.**

Table 11. **Ablation studies on LaSOT.** The default choice for our model is colored in gray .

$$p = \begin{cases} p_{init}, & 0 \leq t < \alpha_1 m, \\ p_{init} \frac{m+t-2\alpha_1 m-\alpha_2 m}{(1-\alpha_1-\alpha_2)m}, & \alpha_1 m \leq t \leq (1-\alpha_2)m, \\ 1.0, & (1-\alpha_2)m < t \leq m, \end{cases} \quad (5)$$

where  $m$  represents the total number of training epochs, and  $t$  is a specific training epoch,  $\alpha_1$  and  $\alpha_2$  are hyper parameters to modulate the training process. Specifically,  $\alpha_1$  controls the duration of warmup process, whereas  $\alpha_2$  determines the length of final finetuning process. The mathematical expectation of  $p$  for each layer is:

$$E(p) = \int_0^m p dt = \left[ \frac{1+p_{init}}{2} + \frac{1-p_{init}}{2}(\alpha_2 - \alpha_1) \right] m. \quad (6)$$

It is worth noting that each layer is optimized fewer times than the total iteration count, according to the mathematical expectation. Through dynamically adjusting the replacement rate  $p$ , we eliminate the requirement of finetuning and accomplish an end-to-end model compression.

### 3.5. Training and Inference

Our CompressTracker is a general framework applicable to various student model architectures. It simplifies the student model optimization by using a straightforward end-to-end training process, without the need for multi-stage methods or complex loss functions. During training, teacher model is frozen and we only optimize student tracker. The total loss for CompressTracker is:

$$L = \lambda_{track} L_{track} + \lambda_{pred} L_{pred} + \lambda_{feat} L_{feat}. \quad (7)$$

After training, the various stages of the student model are combined to create a unified model for the inference phase. Consistent with previous methods [14, 49], a Hanning window penalty is adopted.

## 4. Experiments

### 4.1. Implement Details

Our framework CompressTracker is general and not dependent on a specific transformer structure, hence we select a series of trackers, such as SUTrack [13], ODTrack [52], SeqTrack [12], OSTrack [49] and so on as baseline to verify the effectiveness of our CompressTracker. The training datasets consist of LaSOT [17], TrackingNet [37], GOT-10K [29], and COCO [34], following OSTrack [49] and MixFormerV2 [15]. We set  $\lambda_{track}$  as 1,  $\lambda_{pred}$  as 1, and  $\lambda_{feat}$  as 0.2. The  $p_{init}$  is set as 0.5. We train the CompressTracker with AdamW optimizer [36], with the weight decay as  $10^{-4}$  and the initial learning rate of  $4 \times 10^{-5}$ . The batch size is 128. The total training epochs is 500 with 60K image pairs per epoch and the learning rate is reduced by a factor of 10 after 400 epochs.  $\alpha_1$  and  $\alpha_2$  are set as 0.1. The search and template images are resized to resolutions of  $256 \times 256$  and  $128 \times 128$ . We initialize the CompressTracker with the pretrained parameters. We report the inference speed on a NVIDIA RTX 2080Ti GPU and Intel(R) Xeon(R) Platinum 8268 CPU @ 2.90GHz.

### 4.2. Compress Object Tracker

To validate the effectiveness of CompressTracker, we apply our CompressTracker framework to compress existing transformer-based tracking models, demonstrating its ability to balance performance and inference speed. We carefully




Figure 3. Ablation study on training strategy.

select model configurations that offer an optimal trade-off between accuracy and efficiency. The experimental results are presented in Table 1. Our method achieves significant acceleration while preserving tracking performance. For instance, when compressing SUTrack (73.2 AUC on LaSOT), our CompressTracker maintains 99% performance (72.2 AUC on LaSOT) while achieving a 2.42× speedup. Furthermore, compared to existing lightweight tracking models, CompressTracker consistently outperforms them in both speed and accuracy, highlighting its superiority and strength.

### 4.3. Generalization Verification

In the previous section, we demonstrate the effectiveness of our CompressTracker. To further validate its generalization ability, we conduct additional experiments across diverse settings. The results, shown in Table 2, confirm that our method maintains strong performance across different model architectures, stage numbers, input resolutions, and parameter scales. We also show the results with different structure of student in Supplementary Materials. Moreover, CompressTracker is highly efficient in training, requiring only 20 hours for CompressTracker-OSTrack-4 using 8 NVIDIA RTX 3090 GPUs, in contrast to the 120 hours needed for MixFormerV2-S. This training efficiency further highlights the strong generalization capability of our method, as it can rapidly adapt to various configurations while maintaining high accuracy. Based on the extensive experiment results, our CompressTracker has both highly effectiveness and strong generalization. CompressTracker maximizes performance retention while significantly improving inference speed and is applicable to any *teacher model*, any *resolution*, any *stage number*, any *teacher model size*, and any *student architectures*, which previous methods cannot achieve.

### 4.4. Ablation Study

In this section, we conduct a series of ablation studies on LaSOT to explore the factors contributing to the effectiveness of our CompressTracker. Unless otherwise specified, the teacher model is OSTrack, and the student model has 4 encoder layers. The student model is trained for 300 epochs and  $p_{init}$  is set as 0.5. Please see Supplementary Materials




Figure 4. Ablation study on different replacement probability.

for more analysis.

**Backbone Initialization.** We initialize the backbone of student model with different parameters and only train the student model with groundtruth supervision. The results are shown in Table 3. It can be observed that utilizing the knowledge from teacher model is crucial. Moreover, initializing with skipped layers (#3) yields slightly better performance than continuous layers. This suggests that initialization with skipped layers leads to improved representation similarity.

**Decoder Initialization and Optimization.** We investigate the influence of decoder’s initialization and optimization on the accuracy of student tracker in Table 4. Initializing the decoder with parameters from the teacher model (#2) results in an improvement of approximately 0.3% compared to a decoder initialized randomly (#1), which underscores the benefits of transferring knowledge from the teacher model to enhance the accuracy of the student model’s decoder. Furthermore, making the decoder trainable leads to an additional improvement of 0.2%.

**Stage Division.** Our stage division strategy divides the teacher model into the several stages, and we explore the stage division strategy in Table 5. We design two kinds of division strategy: even and uneven, For the even division, we evenly split the teacher model’s 12 layers into 4 stages, with each stage comprising 3 layers. For uneven division, we follow the design manner in [22, 35] and divide the 12 layers at a ratio of 1:1:3:1. Consequently, the number of layers in each stage of the teacher model is 2, 2, 6, and 2, respectively. The performance of the two approaches is comparable, leading us to select the equal division strategy for simplicity.

**Replacement Probability.** We investigate the impact of replacement probability on the accuracy of student model in Figure 4. We maintain a constant replacement probability instead of implementing the progressive replacement strategy and train the student model with 300 epochs and 30 extra finetuning epochs. It can be observed from Figure 4 that performance is adversely affected when the replacement probability is set either too high or too low. Optimal results are achieved when the replacement probability is within the range of 0.5 to 0.7. Specifically, a too low probability leads

#	Prediction Guidance	Feature Mimicking	Replacement Training	AUC
1				62.8
2	✓			63.5
3		✓		63.3
4			✓	63.7
5	✓		✓	64.1
6		✓	✓	64.5
7	✓	✓		64.3
<b>8</b>	✓	✓	✓	<b>65.2</b>

Table 12. **Ablation study** on LaSOT about supervision of student model. The default choice for our model is colored in gray .

to inadequate training, whereas a too high probability may result in the insufficient interaction between teacher model and student tracker. Thus, we set the  $p_{init}$  as 0.5 based on the experiment result.

**Analysis on Supervision.** We conduct a series of experiments to comprehensively analyze the supervision effects on the student model and to verify the effectiveness of our proposed training strategy. Results are presented in Table 12. Our proposed replacement training approach (#4) improves by 0.9 % AUC compared to singly training student model on groundtruth (#1), which demonstrates that the replacement training enhances the similarity between teacher and student models. Besides, prediction guidance (#5) and feature mimicking (#8) further boost the performance, indicating the effectiveness of the two strategies. Compared to only training on groundtruth (#1), our proposed replacement training, prediction guidance and feature mimicking collectively assist student model in more closely mimicking the teacher model, resulting in a total increase of 2.4% AUC.

To further explore the generalization ability of our proposed training strategy, we compare the performance of models with different layer numbers and training settings, as illustrated in Figure 3. 'Naive Training' denotes that the student model is trained without teacher supervision and replacement training. 'Distill Training' represents that the student model is trained only with teacher supervision. 'CompressTracker' refers to the same training setting in Table 12 #8. It can be observed that as the number of layers increases, there is a corresponding improvement in accuracy. Our CompressTracker shows a noticeable performance boost due to our proposed training strategy, which verifies the effectiveness and generalization ability of our framework.

**Replacement Training.** To evaluate the efficiency and effectiveness of our replacement training strategy, we conduct experiments presented in Table 6. 'Random' denotes our replacement training, and 'Decouple-300' represents stage-by-stage decoupling. Result of # 1 aligns with our replacement training with 300 training epochs, while in # 2, we apply decoupled training, sequentially training and freezing each stage for 75 epochs, followed by 30 epochs of fine-

tuning. The 'Decouple-300' (# 2) approach achieves 64.6% AUC on LaSOT with the same training epochs, marginally lower by 0.6% AUC than our replacement training strategy (# 1). The 'Decouple-300' approach (# 2) requires a complex, multi-stage training along with supplementary fine-tuning, which may suffer from suboptimal outcomes at a specific training process. However, our CompressTracker operates on an end-to-end, single-step basis, and can avoid the suboptimal performance issue through its unified training manner, which validates the superiority of our replacement training.

**Progressive Replacement.** In Table 7, we illustrate the impact of progressive replacement strategy. The first row (# 1) corresponds to the same setting of CompressTracker, while in the second row (# 2) we fix the sampling probability as 0.5 and the student model is trained with 300 epochs followed by 30 finetuning epochs. The absence of progressive replacement leads to a performance degradation of 0.4% AUC, thereby highlighting the efficacy of our progressive replacement approach.

**Training Epochs.** Based on the analysis in Section 3.4, the optimization steps for each layer are lower than total training steps. Thus, to ensure adequate training of each stage, we increase the training epochs from 300 to 500, and show the result in Table 8. Extending the training epochs ensures that student models receive comprehensive training, leading to improved accuracy.

**Training Time.** We compare the training time of our CompressTracker-4 with 500 training epochs, OTrack, and MixFormerV2-S in Table 9. The training time is recorded on 8 NVIDIA RTX 3090 GPUs. Although our CompressTracker requires a longer training time compared to the OTrack, the increased computational overhead remains within acceptable limits. Moreover, MixFormerV2-S is trained on 8 Nvidia RTX8000 GPUs, and we estimate this will take roughly 80 hours on 8 NVIDIA RTX 3090 GPUs based on the relative computational capabilities of these GPUs. The training time of our CompressTracker-4 is significantly less than that of MixFormerV2-S, which validate the efficiency and effectiveness of our framework.

## 5. Conclusion

In this paper, we propose a general compression framework, CompressTracker, for visual object tracking. We propose a novel stage division strategy to separate the structural dependencies between the student and teacher models. We propose the replacement training to enhance student's ability to emulate the teacher model. We further introduce the prediction guidance and stage-wise feature mimicking to improve performance. Extensive experiments verify the effectiveness and generalization ability of our CompressTracker. Our CompressTracker is capable of accelerating tracking models while preserving performance to the greatest extent.

## 6. Acknowledgement

This work was supported by National Natural Science Foundation of China (No.62072112).

## References

- [1] Yifan Bai, Zeyang Zhao, Yihong Gong, and Xing Wei. Ar-trackv2: Prompting autoregressive tracker where to look and how to describe. *arXiv preprint arXiv:2312.17133*, 2023. 3
- [2] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *Computer Vision—ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8-10 and 15-16, 2016, Proceedings, Part II 14*, pages 850–865. Springer, 2016. 3
- [3] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6182–6191, 2019. 4
- [4] Philippe Blatter, Menelaos Kanakis, Martin Danelljan, and Luc Van Gool. Efficient visual tracking with exemplar transformers. In *Proceedings of the IEEE/CVF Winter conference on applications of computer vision*, pages 1571–1581, 2023. 1, 3, 4
- [5] David S Bolme, J Ross Beveridge, Bruce A Draper, and Yui Man Lui. Visual object tracking using adaptive correlation filters. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 2544–2550. IEEE, 2010. 3
- [6] Vasyl Borsuk, Roman Vei, Orest Kupyn, Tetiana Martyniuk, Igor Krashenyi, and Jiří Matas. Fear: Fast, efficient, accurate and robust visual tracker. In *European Conference on Computer Vision*, pages 644–663. Springer, 2022. 1, 3, 4
- [7] Arnav Chavan, Zhiqiang Shen, Zhuang Liu, Zechun Liu, Kwang-Ting Cheng, and Eric P Xing. Vision transformer slimming: Multi-dimension searching in continuous optimization space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4931–4941, 2022. 4
- [8] 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 *European Conference on Computer Vision*, pages 375–392. Springer, 2022. 3, 4
- [9] Minghao Chen, Houwen Peng, Jianlong Fu, and Haibin Ling. Autoformer: Searching transformers for visual recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12270–12280, 2021. 4
- [10] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8126–8135, 2021. 3
- [11] Xin Chen, Ben Kang, Dong Wang, Dongdong Li, and Huchuan Lu. Efficient visual tracking via hierarchical cross-attention transformer. In *European Conference on Computer Vision*, pages 461–477. Springer, 2022. 1, 3, 4
- [12] Xin Chen, Houwen Peng, Dong Wang, Huchuan Lu, and Han Hu. Seqtrack: Sequence to sequence learning for visual object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14572–14581, 2023. 3, 4, 5, 6
- [13] Xin Chen, Ben Kang, Wanting Geng, Jiawen Zhu, Yi Liu, Dong Wang, and Huchuan Lu. Sutrack: Towards simple and unified single object tracking. *arXiv preprint arXiv:2412.19138*, 2024. 4, 5, 6
- [14] Yutao Cui, Cheng Jiang, Limin Wang, and Gangshan Wu. Mixformer: End-to-end tracking with iterative mixed attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13608–13618, 2022. 3, 4, 6
- [15] Yutao Cui, Tianhui Song, Gangshan Wu, and Limin Wang. Mixformerv2: Efficient fully transformer tracking. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 3, 4, 6
- [16] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Atom: Accurate tracking by overlap maximization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4660–4669, 2019. 3
- [17] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5374–5383, 2019. 6
- [18] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018. 3
- [19] Shenyuan Gao, Chunlun Zhou, and Jun Zhang. Generalized relation modeling for transformer tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18686–18695, 2023. 3
- [20] Chengyue Gong and Dilin Wang. Nasvit: Neural architecture search for efficient vision transformers with gradient conflict-aware supernet training. *ICLR Proceedings 2022*, 2022. 4
- [21] Goutam Yelluru Gopal and Maria A Amer. Separable self and mixed attention transformers for efficient object tracking. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6708–6717, 2024. 1, 4
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [23] João F Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista. High-speed tracking with kernelized correlation filters. *IEEE transactions on pattern analysis and machine intelligence*, 37(3):583–596, 2014. 3
- [24] Jack Hong, Shilin Yan, Zehao Xiao, Jiayin Cai, Xiaolong Jiang, Yao Hu, and Henghui Ding. Progressive scaling visual object tracking. *arXiv preprint arXiv:2505.19990*, 2025. 3
- [25] Lingyi Hong, Wenchao Chen, Zhongying Liu, Wei Zhang, Pinxue Guo, Zhaoyu Chen, and Wenqiang Zhang. Lvos:

- A benchmark for long-term video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13480–13492, 2023. 3
- [26] Lingyi Hong, Wei Zhang, Shuyong Gao, Hong Lu, and Wen-Qiang Zhang. Simulflow: Simultaneously extracting feature and identifying target for unsupervised video object segmentation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 7481–7490, 2023. 3
- [27] Lingyi Hong, Zhongying Liu, Wenchao Chen, Chenzhi Tan, Yuang Feng, Xinyu Zhou, Pinxue Guo, Jinglun Li, Zhaoyu Chen, Shuyong Gao, et al. Lvos: A benchmark for large-scale long-term video object segmentation. *arXiv preprint arXiv:2404.19326*, 2024. 3
- [28] Lingyi Hong, Shilin Yan, Renrui Zhang, Wanyun Li, Xinyu Zhou, Pinxue Guo, Kaixun Jiang, Yiting Chen, Jinglun Li, Zhaoyu Chen, et al. Onetracker: Unifying visual object tracking with foundation models and efficient tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19079–19091, 2024. 3
- [29] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE transactions on pattern analysis and machine intelligence*, 43(5):1562–1577, 2019. 6
- [30] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*, 2019. 3
- [31] Ben Kang, Xin Chen, Dong Wang, Houwen Peng, and Huchuan Lu. Exploring lightweight hierarchical vision transformers for efficient visual tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9612–9621, 2023. 1, 3, 4
- [32] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8971–8980, 2018. 3
- [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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4282–4291, 2019. 3
- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 6
- [35] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022. 7
- [36] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [37] 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 *Proceedings of the European conference on computer vision (ECCV)*, pages 300–317, 2018. 6
- [38] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. *Advances in neural information processing systems*, 34:13937–13949, 2021. 4
- [39] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019. 3
- [40] Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Q-bert: Hessian based ultra low precision quantization of bert. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8815–8821, 2020. 3
- [41] Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. Patient knowledge distillation for bert model compression. *arXiv preprint arXiv:1908.09355*, 2019. 3
- [42] Xing Wei, Yifan Bai, Yongchao Zheng, Dahu Shi, and Yihong Gong. Autoregressive visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9697–9706, 2023. 3
- [43] Canwen Xu, Wangchunshu Zhou, Tao Ge, Furu Wei, and Ming Zhou. Bert-of-theseus: Compressing bert by progressive module replacing. *arXiv preprint arXiv:2002.02925*, 2020. 3
- [44] Yinda Xu, Zeyu Wang, Zuoxin Li, Ye Yuan, and Gang Yu. Siamfc++: Towards robust and accurate visual tracking with target estimation guidelines. In *Proceedings of the AAAI conference on artificial intelligence*, pages 12549–12556, 2020. 4
- [45] Yifan Xu, Zhijie Zhang, Mengdan Zhang, Kekai Sheng, Ke Li, Weiming Dong, Liqing Zhang, Changsheng Xu, and Xing Sun. Evo-vit: Slow-fast token evolution for dynamic vision transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2964–2972, 2022. 4
- [46] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10448–10457, 2021. 3, 4
- [47] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15189, 2021. 1, 3, 4
- [48] Zhendong Yang, Zhe Li, Ailing Zeng, Zexian Li, Chun Yuan, and Yu Li. Vitkd: Practical guidelines for vit feature knowledge distillation. *arXiv preprint arXiv:2209.02432*, 2022. 4
- [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 *European conference on computer vision*, pages 341–357. Springer, 2022. 1, 3, 4, 5, 6

- [50] Jinnian Zhang, Houwen Peng, Kan Wu, Mengchen Liu, Bin Xiao, Jianlong Fu, and Lu Yuan. Minivit: Compressing vision transformers with weight multiplexing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12145–12154, 2022. [4](#)
- [51] Zhipeng Zhang, Houwen Peng, Jianlong Fu, Bing Li, and Weiming Hu. Ocean: Object-aware anchor-free tracking. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pages 771–787. Springer, 2020. [3](#)
- [52] Yaozong Zheng, Bineng Zhong, Qihua Liang, Zhiyi Mo, Shengping Zhang, and Xianxian Li. Odtrack: Online dense temporal token learning for visual tracking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7588–7596, 2024. [4](#), [5](#), [6](#)
- [53] Xinyu Zhou, Pinxue Guo, Lingyi Hong, Jinglun Li, Wei Zhang, Weifeng Ge, and Wenqiang Zhang. Reading relevant feature from global representation memory for visual object tracking. *Advances in Neural Information Processing Systems*, 36:10814–10827, 2023. [3](#)
- [54] Xinyu Zhou, Jinglun Li, Lingyi Hong, Kaixun Jiang, Pinxue Guo, Weifeng Ge, and Wenqiang Zhang. Detrack: In-model latent denoising learning for visual object tracking. *Advances in Neural Information Processing Systems*, 37:90579–90599, 2024. [3](#)