

# Generalized Relation Modeling for Transformer Tracking

Shenyuan Gao<sup>1</sup> Chunlun Zhou<sup>2</sup> Jun Zhang<sup>1</sup>

<sup>1</sup>The Hong Kong University of Science and Technology <sup>2</sup>Wormpex AI Research

sygao@connect.ust.hk czhou002@e.ntu.edu.sg eejzhang@ust.hk

## Abstract

Compared with previous two-stream trackers, the recent one-stream tracking pipeline, which allows earlier interaction between the template and search region, has achieved a remarkable performance gain. However, existing one-stream trackers always let the template interact with all parts inside the search region throughout all the encoder layers. This could potentially lead to target-background confusion when the extracted feature representations are not sufficiently discriminative. To alleviate this issue, we propose a generalized relation modeling method based on adaptive token division. The proposed method is a generalized formulation of attention-based relation modeling for Transformer tracking, which inherits the merits of both previous two-stream and one-stream pipelines whilst enabling more flexible relation modeling by selecting appropriate search tokens to interact with template tokens. An attention masking strategy and the Gumbel-Softmax technique are introduced to facilitate the parallel computation and end-to-end learning of the token division module. Extensive experiments show that our method is superior to the two-stream and one-stream pipelines and achieves state-of-the-art performance on six challenging benchmarks with a real-time running speed. Code and models are publicly available at <https://github.com/Little-Podi/GRM>.

## 1. Introduction

Given the target bounding box in the initial frame of a video, visual tracking [18] aims to localize the target in successive frames. Over the past few years, two-stream trackers [1, 21, 22, 49], which extract features of the template and search region separately and then model cross-relations of the template and search region in a sequential fashion, have emerged as a dominant tracking paradigm and made a significant progress. Following this two-stream pipeline, several Transformer-based trackers [4, 11, 38] utilize parallel self-attention blocks to enhance the extracted features by modeling global self-relations within each image as illustrated in Fig. 1(a). Recently, leveraging

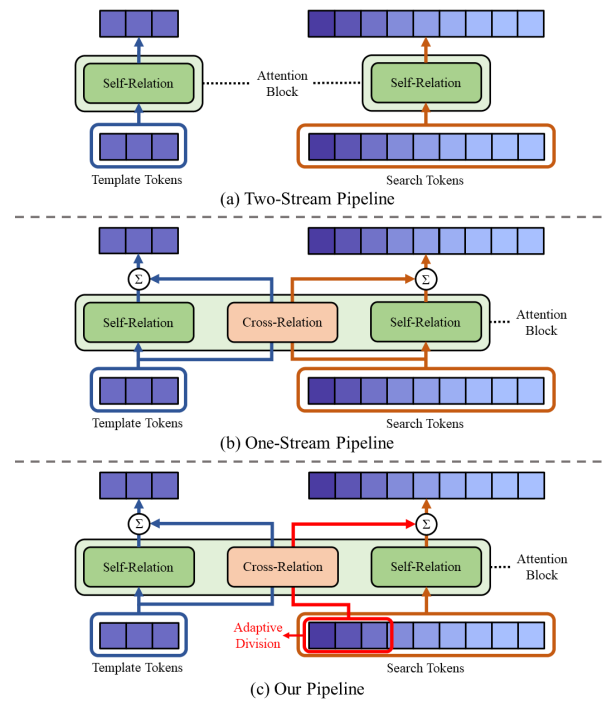


Figure 1. Comparison of different relation modeling pipelines of Transformer-based trackers. The two-stream pipeline uses parallel self-attention blocks to model relations within each set of tokens (template or search tokens). The one-stream pipeline integrates the cross-relation modeling between two sets of tokens and self-relation modeling within each set of tokens via a unified attention block. In contrast, our proposed pipeline performs an adaptive division of the search tokens, which can degenerate to the two-stream form if no search token is selected to interact with template tokens and to the one-stream form if all the search tokens are selected for cross-relation modeling.

the flexibility of the attention mechanism, the one-stream pipeline [3, 5, 44] is proposed to jointly extract features and model relations, achieving promising performance. By conducting self-attention among all concatenated tokens, both cross-relation modeling and self-relation modeling can be performed simultaneously as illustrated in Fig. 1(b).

It is demonstrated in [3, 5, 39, 44] that letting the search region interact with the template as early as possible is ben-

official to target-specific feature generation. However, there is no evidence suggesting that all parts inside the search region should always be forced to interact with the template. Actually, due to the cropping strategy [1], there is a large proportion of background inside the search region, where distractors with similar appearance to the target may exist. This would lead to undesired cross-relations between the template and search region as the highly discriminative representations have not been extracted in some early layers.

Although the attention mechanism can inherently weaken improper cross-relations, applying global cross-relation modeling to all layers may still be more or less disruptive. On the one hand, for the search tokens outside the target region, if undesired cross-relations are modeled between the template and distractors, the aggregated features of the distractors may contain the target features from the template, which could cause confusion for precisely identifying the actual target in the search region. On the other hand, for the template tokens, their quality could also be degraded by undesired cross-relations during the iterative update since certain features from the background or even distractors could be aggregated into these tokens. These situations could weaken the target-background discrimination capability of the one-stream pipeline.

Intuitively, only a portion of search tokens, *e.g.*, tokens belonging to the target, are suitable for cross-relation modeling when the feature representations are not perfect for target-background discrimination. In some cases, the two-stream relation modeling pipeline could even be better if the feature representations of both the template and search region are imperfect to model cross-relations. The potential limitations of the one-stream pipeline motivates us to ponder: is it really optimal for the template to interact with all parts inside the search region through all encoder layers in the one-stream pipeline?

In this paper, we answer this question by proposing GRM, a generalized relation modeling method that can adaptively select the appropriate search tokens to interact with the template. To be specific, we classify the template and search tokens as three categories. The template tokens form one category while the search tokens are divided into another two categories. Instead of modeling relations within all the tokens as the one-stream pipeline, we restrict the interaction among the three token categories. Only the search tokens that are suitable for cross-relation modeling will interact with the template tokens, whilst the interaction between the remaining search tokens and the template tokens is blocked. With proper divisions, the two-stream pipeline and one-stream pipeline become two degenerated forms of our relation modeling method as discussed in Sec. 3.2. Consequently, our method is a generalized formulation of attention-based relation modeling for Transformer tracking, which embraces the advantages of both previous

pipelines while being more flexible.

The search token division is performed by a lightweight prediction module, which can adaptively determine which search tokens are suitable for cross-relation modeling based on the input tokens. To accomplish this objective, there are two obstacles to overcome. First, the separate relation modeling for different token categories makes it hard for parallel computation. Second, the discrete token categorization is non-differentiable, thus impeding the end-to-end learning of the token division module. To facilitate parallel computation, we adopt an attention masking strategy to unify the individual attention operations into a single one. Additionally, we introduce the Gumbel-softmax technique [17] to make the discrete token categorization differentiable. Consequently, the search token division module can be implicitly optimized in an end-to-end manner, which promotes its adaptability to deal with different situations.

In summary, our main contributions are three-fold:

- We present a generalized formulation of relation modeling for Transformer trackers, which divides the input tokens into three categories and enables more flexible interaction between the template and search region.
- To realize the generalized relation modeling, we devise a token division module to adaptively classify the input tokens. An attention masking strategy and the Gumbel-Softmax technique are introduced to facilitate the parallel computation and end-to-end learning of the proposed module.
- We conduct extensive experiments and analyses to validate the efficacy of our method. The proposed GRM exhibits outstanding results on six challenging visual tracking benchmarks.

## 2. Related Works

### 2.1. Visual Tracking Paradigms

Over the past few years, Siamese trackers [1, 6, 22] have gained much popularity. Typically, they adopt a two-stream pipeline to separately extract the features of the template and search region. Cross-relations between the two streams are modeled by additional correlation modules. To exploit the power of highly discriminative features, most Siamese trackers [2, 21, 48, 49] use deep neural networks as the backbone, *e.g.*, ResNet-50 [14]. Recently, the Transformer [37] architecture has achieved promising results in visual tracking and has become the de-facto choice for many high-performance trackers [4, 11, 23, 28, 38, 40, 41].

To further enhance the feature interaction, several attempts [12, 39, 46] have investigated cross-relation modeling inside the backbone. Recently, another thread of progress [23, 41] concatenates the template and search tokens to conduct cross-relation modeling and self-relation

modeling jointly. Inspired by these explorations, more recent trackers [3, 5, 44] adopt a one-stream pipeline to jointly extract the features and model the relations of both the template and search region by the self-attention mechanism. Based on this pipeline, they can utilize advanced pretrained models, *e.g.*, MAE [13], instead of randomly initialized correlation modules for cross-relation modeling, thereby yielding a remarkable performance gain. Akin to most previous two-stream methods, however, these one-stream methods also treat the search region as a whole, which means that the template always interacts with all parts inside the search region. This may unfavorably cause target-background confusion when the feature representations are imperfect for cross-relation modeling. To mitigate such confusion, we devise a lightweight module to predict which search tokens are appropriate for the cross-relation modeling with the template in each encoder layer. We also show that both the two-stream and one-stream pipelines become degenerated forms of our method.

## 2.2. Trackers with Dynamic Designs

There are several trackers leveraging dynamic designs at different stages for various purposes. Among them, LightTrack [42] and AutoMatch [47] use neural architecture search techniques during the training stage. Particularly, LightTrack aims to find a more efficient structure of the network backbone and prediction head, whilst AutoMatch aims to automatically design a more powerful correlation module. Both the network architectures of these trackers are frozen after the training stage. On the contrary, our method introduces a dynamic architecture design that can conduct flexible adaptations during the inference stage.

Benefiting from the flexibility of the attention mechanism, several efforts [20, 26, 34, 45] have been made to design different token organization methods to accelerate the inference speed of vision Transformer. To make it well-suited to the task of visual tracking, SparseTT [10] and OSTRack [44] utilize the attention weights that are free to borrow from its self-attention calculation. Concretely, SparseTT only aggregates the features from the most similar tokens to improve the discrimination ability, whilst OSTRack eliminates the background tokens to boost efficiency. However, these operations do not have learnable parameters and need to be hand-crafted. In addition, a large proportion of search tokens are still consistently involved in the cross-relation modeling with the template, which may lead to suboptimal feature aggregation since undesired interaction are implicated when the feature representations are not sufficiently discriminative. Differently, we propose an end-to-end learnable token division for Transformer visual tracking, which can adaptively select appropriate search tokens for cross-relation modeling with the template.

## 3. Method

### 3.1. Preliminary

We first revisit the recent one-stream pipeline [3, 5, 44], which achieves remarkable tracking performance. The input of the one-stream pipeline is a pair of images, namely template  $z \in \mathbb{R}^{H_z \times W_z \times 3}$  and search region  $x \in \mathbb{R}^{H_x \times W_x \times 3}$ . They are first divided into  $N_z$  and  $N_x$  non-overlapping image patches of resolution  $P \times P$  respectively, where  $N_z = H_z W_z / P^2$  and  $N_x = H_x W_x / P^2$ . Then, a linear projection is applied to these image patches to generate template patch embeddings  $\mathbf{E}_z \in \mathbb{R}^{N_z \times C}$  and search patch embeddings  $\mathbf{E}_x \in \mathbb{R}^{N_x \times C}$ , where  $C$  is the embedding dimension. Next, two learnable position embeddings are added to  $\mathbf{E}_z$  and  $\mathbf{E}_x$  respectively to embed the spatial information. With some abuse of notation, we still refer to the two updated embeddings as  $\mathbf{E}_z$  and  $\mathbf{E}_x$ , which are the initial token embeddings of the template and search region, respectively. Afterward, all these tokens are concatenated as a sequence with a length of  $N_z + N_x$  and fed to an encoder. Each encoder layer updates the input tokens via a multi-head attention (MHA) block and a feed-forward network (FFN). Formally, the operations of the  $l$ -th encoder layer can be expressed as:

$$\begin{aligned} \mathbf{q} = \mathbf{k} = \mathbf{v} &= [\mathbf{E}_z^l; \mathbf{E}_x^l], \\ [\mathbf{E}'_z^l; \mathbf{E}'_x^l] &= [\mathbf{E}_z^l; \mathbf{E}_x^l] + \text{MHA}(\mathbf{q}, \mathbf{k}, \mathbf{v}), \\ [\mathbf{E}_z^{l+1}; \mathbf{E}_x^{l+1}] &= [\mathbf{E}'_z^l; \mathbf{E}'_x^l] + \text{FFN}([\mathbf{E}'_z^l; \mathbf{E}'_x^l]), \end{aligned} \quad (1)$$

where  $\mathbf{E}_z^l$  and  $\mathbf{E}_x^l$  are the input tokens of the  $l$ -th encoder layer, and  $[\cdot; \cdot]$  denotes the concatenation operation. We use  $\mathbf{q}$ ,  $\mathbf{k}$  and  $\mathbf{v}$  to represent the queries, keys and values fed to the multi-head attention block. Since the template and search tokens are jointly processed by the multi-head attention block, the cross-relation and self-relation modeling are seamlessly integrated in each encoder layer. Finally, the output search tokens from the last encoder layer are decoupled from the template tokens and re-shaped to a 2D feature map according to their original spatial positions. The feature map is then taken as the input of to a convolutional head for target bounding box prediction.

### 3.2. Generalized Relation Modeling

The one-stream pipeline allows free interaction between template tokens and search tokens via the attention mechanism in every encoder layer, which facilitates target-specific feature extraction. However, the large proportion of background inside the search region could lead to undesired cross-relations as the highly discriminative representations have not been extracted in some early layers. Although the attention mechanism can inherently weaken improper cross-relations, they may still cause adverse effects. On the one hand, a certain amount of information from the

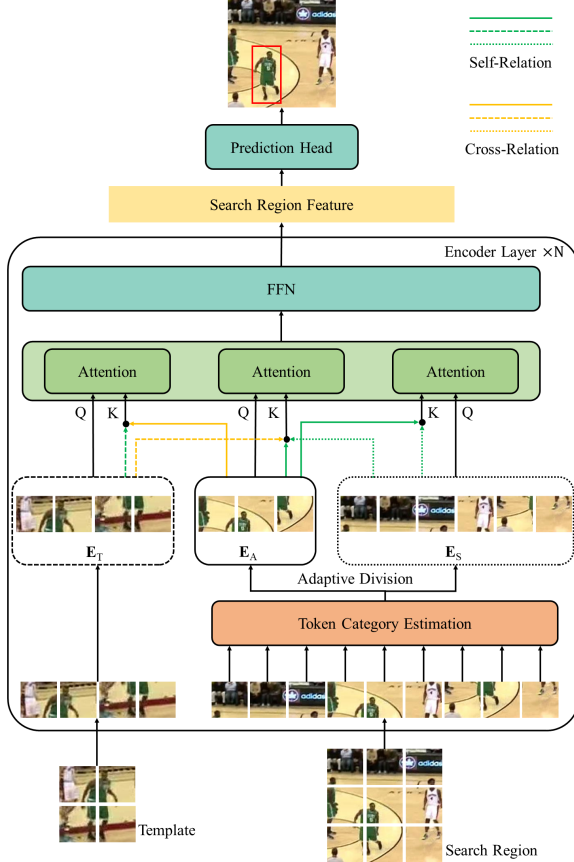


Figure 2. Tracking framework with our generalized relation modeling method. The attention operations for the three token categories are simultaneously conducted with the proposed attention masking strategy. We only illustrate the queries and keys for the attention operations of different token categories. The values are identical to the corresponding keys and are omitted for clarity.

template is aggregated into the search region that belong to background or distractors, which could increase the difficulty of identifying the target. On the other hand, the representation of the template tokens could be distracted by the information from the inappropriate search tokens, which degrades the quality of the template during the iterative update. Therefore, it may not be optimal for the template to always interact with all parts inside the search region in every encoder layer.

To mitigate the aforementioned issues, a straightforward solution is to divide the input tokens  $\mathbf{E}$  into two categories, denoted by  $\mathbf{E}_T$  and  $\mathbf{E}_S$ . The first category  $\mathbf{E}_T$  contains all the template tokens and a portion of search tokens for which the relation modeling with the template tokens are necessary, whilst the second category  $\mathbf{E}_S$  contains the remaining search tokens. Then, the attention mechanism can be applied to model relations within tokens of each category. This token division can help prevent the undesired interaction between the template tokens and inappropriate search tokens. However, if we directly apply this binary

division, there is a risk that the relation modeling between some relevant search tokens would be blocked. For example, when the search tokens of one object are separated into two sets, the interaction between these two sets would not be allowed, which is adverse to feature aggregation.

To overcome the downside of this binary division, we introduce another token category, denoted by  $\mathbf{E}_A$ , to ensure sufficient interaction between relevant search tokens while preventing the undesired interaction between template tokens and inappropriate search tokens. We divide the search tokens into two sets, one is assigned to category  $\mathbf{E}_A$  and the other is assigned to category  $\mathbf{E}_S$ . We fix all the template tokens as the category  $\mathbf{E}_T$  (i.e.,  $\mathbf{E}_T = \mathbf{E}_z$ ). The relation modeling rules for the three categories are defined as:

1. For category  $\mathbf{E}_T$ , its tokens can aggregate information from the tokens in  $\mathbf{E}_T$  and  $\mathbf{E}_A$ :

$$\begin{aligned} \mathbf{q} &= \mathbf{E}_T^l = \mathbf{E}_z^l, \\ \mathbf{k} &= \mathbf{v} = [\mathbf{E}_T^l; \mathbf{E}_A^l] = [\mathbf{E}_z^l; \mathbf{E}_A^l], \\ \mathbf{E}'_T &= \mathbf{E}_T^l + \text{MHA}(\mathbf{q}, \mathbf{k}, \mathbf{v}). \end{aligned} \quad (2)$$

2. For category  $\mathbf{E}_S$ , its tokens can aggregate information from the tokens in  $\mathbf{E}_S$  and  $\mathbf{E}_A$ :

$$\begin{aligned} \mathbf{q} &= \mathbf{E}_S^l, \\ \mathbf{k} &= \mathbf{v} = [\mathbf{E}_S^l; \mathbf{E}_A^l] = \mathbf{E}_x^l, \\ \mathbf{E}'_S &= \mathbf{E}_S^l + \text{MHA}(\mathbf{q}, \mathbf{k}, \mathbf{v}). \end{aligned} \quad (3)$$

3. For category  $\mathbf{E}_A$ , its tokens can aggregate information from all tokens:

$$\begin{aligned} \mathbf{q} &= \mathbf{E}_A^l, \\ \mathbf{k} &= \mathbf{v} = [\mathbf{E}_T^l; \mathbf{E}_S^l; \mathbf{E}_A^l] = [\mathbf{E}_z^l; \mathbf{E}_x^l], \\ \mathbf{E}'_A &= \mathbf{E}_A^l + \text{MHA}(\mathbf{q}, \mathbf{k}, \mathbf{v}). \end{aligned} \quad (4)$$

With  $\mathbf{E}'_T$ ,  $\mathbf{E}'_S$  and  $\mathbf{E}'_A$ , the output tokens of the  $l$ -th encoder are obtained by:

$$\begin{aligned} \mathbf{E}'_z &= \mathbf{E}'_T, \\ \mathbf{E}'_x &= [\mathbf{E}'_S; \mathbf{E}'_A], \\ [\mathbf{E}_z^{l+1}; \mathbf{E}_x^{l+1}] &= [\mathbf{E}'_z; \mathbf{E}'_x] + \text{FFN}([\mathbf{E}'_z; \mathbf{E}'_x]). \end{aligned} \quad (5)$$

Fig. 2 illustrates our generalized relation modeling method for the three token categories in each encoder layer. The template tokens in  $\mathbf{E}_T$  and search tokens in  $\mathbf{E}_A$  can interact with each other via the attention operations defined in Eq. (2) and Eq. (4). The interaction between the template tokens and the remaining search tokens in  $\mathbf{E}_S$  are blocked. All the search tokens can fully interact with each other via the attention operations defined in Eq. (3) and Eq. (4).

Our proposed relation modeling method is a generalized formulation of previous two-stream and one-stream pipelines. Concretely, for the two-stream form, our method degenerates to it when all the search tokens are assigned to  $\mathbf{E}_S$ . The template tokens interact with themselves via an attention operation (Eq. (2)) and the search tokens interact with themselves via another attention operation (Eq. (3)), forming two-stream relation modeling pipeline. As the category  $\mathbf{E}_A$  is empty, no cross-relation modeling is performed. For the one-stream form, our method degenerates to it when all the search tokens are assigned to  $\mathbf{E}_A$ . In this case, the category  $\mathbf{E}_S$  contains no search tokens and the operations in Eq. (2) and Eq. (4) can be combined as a single attention operation. Thus, all the template tokens and search tokens can interact with each other, forming one-stream relation modeling pipeline. Our generalized design can provide higher flexibility to handle different situations in different layers. As shown in Sec. 4.4, our relation modeling method can adaptively degenerate to the two-stream and one-stream forms when necessary.

### 3.3. Adaptive Token Division

To realize our generalized relation modeling method, we need a way to dynamically divide search tokens into either category  $\mathbf{E}_S$  or category  $\mathbf{E}_A$ . To achieve this, we devise a learnable prediction module for each encoder layer. In particular, to provide target-related cues for search tokens, we generate a target-aware representation by aggregating all the template tokens with global maximum pooling. The target-aware representation is then concatenated with each search token. Afterward, we send the concatenation to a light-weight multi-layer perceptron (MLP) for predicting the probabilities of search tokens belonging to category  $\mathbf{E}_S$  and category  $\mathbf{E}_A$ :

$$\pi = \text{Softmax}(\text{MLP}([\text{MaxPool}(\mathbf{E}_z); \mathbf{E}_x])) \in \mathbb{R}^{N_x \times 2}, \quad (6)$$

where  $\pi_{i,0}$  and  $\pi_{i,1}$  denote the probabilities of the  $i$ -th search token being assigned to  $\mathbf{E}_S$  and  $\mathbf{E}_A$ , respectively. With the prediction  $\pi$ , we determine the category of the  $i$ -th search token by selecting the category with a higher probability.

Although it seems to be straightforward to apply the prediction module to perform search token division, it is non-trivial for implementation because there are still two obstacles to tackle. That are, how to parallelize the computation for GPU platforms and how to learn the proposed prediction module through end-to-end optimization.

The first obstacle arises from the different relation modeling rules for the three token categories. Generally, the numbers of queries and keys are different for the three categories, making it hard to parallelize the three attention operations in Eq. (2), Eq. (3) and Eq. (4). In practice, we empirically find that the speed will be significantly slowed

down if we conduct the attention operations for the three categories separately. Inspired by DynamicViT [34], we adopt an attention masking strategy to tackle this problem. We first translate probabilities  $\pi$  to the one-hot categorization  $\mathbf{D} \in \{0, 1\}^{N_x \times 2}$ . Then, we define two one-hot tensors,  $\mathbf{D}_z \in \{0, 1\}^{N_z \times 3}$  and  $\mathbf{D}_x \in \{0, 1\}^{N_x \times 3}$ , which represent the categorization of template tokens and that of search tokens, respectively. The first column of  $\mathbf{D}_z$  is filled with ones and the other two columns are set to zeros, indicating all the template tokens are fixed as  $\mathbf{E}_T$ . The one-hot tensor  $\mathbf{D}_x$  is obtained by appending a column of zeros to the front of the binary categorization  $\mathbf{D}$ , meaning no search tokens will be classified as  $\mathbf{E}_T$ . With the categorization of all the tokens  $\widehat{\mathbf{D}} = [\mathbf{D}_z; \mathbf{D}_x] \in \{0, 1\}^{(N_z+N_x) \times 3}$ , we then construct an attention mask  $\mathbf{M} \in \{0, 1\}^{(N_z+N_x) \times (N_z+N_x)}$ , where each value  $M_{i,j}$  is calculated by:

$$\mathbf{M}_{i,j} = \widehat{\mathbf{D}}_{i,0} \widehat{\mathbf{D}}_{j,0} + \widehat{\mathbf{D}}_{i,1} \widehat{\mathbf{D}}_{j,1} + \widehat{\mathbf{D}}_{i,2} (\widehat{\mathbf{D}}_{j,0} + \widehat{\mathbf{D}}_{j,1} + \widehat{\mathbf{D}}_{j,2}). \quad (7)$$

It is easy to verify that  $\mathbf{M}_{i,j}$  indicates whether token  $i$  can aggregate information from token  $j$  according to the relation modeling rules in Eq. (2), Eq. (3) and Eq. (4). By computing the Hadamard products of the attention mask and the attention weight matrix of all input tokens, we can combine the three attention operations into a single one with the same functionality. The introduced attention masking strategy can facilitate computation parallelization on GPUs with a minor increase in computational cost.

Second, the discrete categorization of search tokens is non-differentiable, which impedes the optimization of the prediction module. To overcome this drawback, we apply the Gumbel-Softmax technique [17] to draw samples from a categorical distribution with class probabilities  $\pi$  by a reparameterization trick:

$$\mathbf{D} = \text{One-hot}(\arg \max_i [g_i + \log \pi_i]) \in \{0, 1\}^{N_x \times 2}, \quad (8)$$

where  $g \sim \text{Gumbel}(0, 1)$  and the argmax operation is replaced by the softmax operation as a continuous and differentiable approximation:

$$\widehat{y}_i = \frac{\exp((\log(\pi_i) + g_i)/\tau)}{\sum_j \exp((\log(\pi_j) + g_j)/\tau)}. \quad (9)$$

During the training stage, the discrete categorization  $\mathbf{D}$  sampled by Eq. (8) is used to divide search tokens in the forward pass and gradients are computed from the continuous Gumbel-Softmax approximation in Eq. (9) in the backward pass. We refer readers to [17] for the detailed derivation. Thus, the division of search tokens can be implicitly learned by the supervision from the target localization loss presented in Sec. 3.4 in an end-to-end manner.

### 3.4. Target Prediction Head

Our target prediction head is adopted from [44], which receives a 2D feature map to predict the location of the target. It consists of three convolutional branches which are responsible for center classification, offset regression and size regression, respectively. The center classification branch outputs a score map, where each score represents the confidence of the target center locating at the corresponding position. The offset regression branch is utilized to compensate for the discretization error. The size regression branch predicts the height and width of the target. The position with the highest confidence in the center score map is selected as the target position and the corresponding regressed coordinates are used to compute a bounding box as the final prediction.

During the training stage, the classification branch is supervised by a Gaussian map generated according to the center of the ground truth. For the two regression branches, only the coordinates selected to obtain the final prediction are taken account in the regression loss. The total loss of the prediction head is a weighted sum of a focal loss [24] for bounding box center classification, a GIoU loss [35] and a L1 loss for bounding box coordinate regression:

$$L_{total} = \lambda_{center}L_{focal} + \lambda_{giou}L_{giou} + \lambda_{l1}L_{l1}, \quad (10)$$

where  $\lambda_{center}$ ,  $\lambda_{giou}$  and  $\lambda_{l1}$  are trade-off weights to balance the joint optimization.

## 4. Experiments

### 4.1. Implementation Details

The architecture of our network backbone is identical to the encoder part of ViT-B [8]. It includes 12 sequential encoder layers, each of which comprises a multi-head attention block and a feed-forward network. We adopt the MAE-pretrained model [13] to initialize our network backbone for its fast convergence as in [44]. The search region is  $4^2$  times the target object area and resized to a resolution of  $256 \times 256$  pixels, whilst the template is  $2^2$  times the target object area and resized to  $128 \times 128$  pixels. The cropped images are then down-sampled by a patch embedding projection with a stride of 16. The feature resolution is the same from the first to the last encoder layer. Each branch of the target prediction head comprises 4 stacked Conv-BN-ReLU layers. We use an identical structure for all token division modules in different encoder layers, which is a simple MLP that includes two hidden layers with GELU activation [15] and an output layer. The channel dimensions of the hidden layers and output layer are 384, 192 and 2, respectively.

Our experiments are conducted with NVIDIA GeForce RTX 3090 GPUs. In line with most Transformer-based trackers, we use the training splits of LaSOT [9], TrackingNet [31], GOT-10k [16], and COCO [25] for training

except for the GOT-10k evaluation. The whole network is optimized with the AdamW optimizer [27] for 300 epochs. The initial learning rate is  $4 \times 10^{-5}$  for the parameters with the pretrained weights and  $4 \times 10^{-4}$  for other randomly initialized parameters. It decays by a factor of 10 after 240 epochs of training. As for GOT-10k, the model is trained for 100 epochs and the learning rate decays at epoch 80.

### 4.2. Results and Comparisons

To show the effectiveness and generalization ability of the proposed method, we evaluate and compare our GRM with several state-of-the-art trackers on six benchmarks. The results are summarized in Table 1, Table 2 and Table 3. **GOT-10k.** GOT-10k [16] provides 180 short-term video sequences without publicly available ground truths. To ensure zero overlaps of object classes between training and testing, we strictly follow their one-shot protocol. As shown in Table 1, our tracker improves all metrics by a large margin. Notably, with the aligned model settings, our tracker outperforms the most competitive tracker OTrack [44] by 2.4% in AO and 2.5% in success rate.

**TrackingNet.** TrackingNet [31] is a large-scale short-term tracking benchmark. Our results are reported by the online evaluation server. Table 1 shows that our GRM achieves 84.0% in success score and 83.3% in precision score, overtaking all previously published trackers.

**LaSOT.** LaSOT [9] is a densely annotated large-scale dataset that contains 280 long-term video sequences for public evaluation. From Table 1, we find that our method sets a new state-of-the-art on LaSOT, which demonstrates that our approach is also well-suited to the tracking scenarios when the video sequences are extremely long.

**AViST.** AViST [32] is a recently released benchmark, which comprises 120 sequences in diverse scenarios with adverse visibility, such as bad weather conditions and camouflage. Our tracker surpasses the previous best tracker MixFormer [5], which evidences that our method has a stronger capability in handling these challenging scenarios.

**NfS30.** Need for Speed (NfS) [19] comprises 100 videos with fast-moving objects. As previous works, we evaluate our method on its low frame rate version NfS30. The results in Table 2 show that our method achieves comparable performance with the state-of-the-art trackers, demonstrating the generalization ability of GRM.

**UAV123.** UAV123 [30] contains 123 video sequences captured from a low-altitude unmanned aerial vehicle perspective. It is also a long-term tracking dataset with an average sequence length of 915 frames. The result of our method is shown in Table 2. The top-ranked performance shows the strong ability of our tracker to perform long-term tracking.

**Scalability.** To show our scalability to larger architectures and higher resolutions, we also implement a variant named GRM-L320. With all other configurations fixed, we change

Table 1. State-of-the-art comparison on GOT-10k, TrackingNet, LaSOT and AViS-T. The best two results are shown in red and blue fonts, respectively. We use \* to denote that the results on GOT-10k are obtained following the official one-shot protocol.

Tracker	Source	GOT-10k* [16]			TrackingNet [31]			LaSOT [9]			AViS-T [32]		
		AO	SR <sub>0.5</sub>	SR <sub>0.75</sub>	AUC	P <sub>Norm</sub>	P	AUC	P <sub>Norm</sub>	P	AUC	OP50	OP75
GRM	Ours	<b>73.4</b>	<b>82.9</b>	<b>70.4</b>	<b>84.0</b>	<b>88.7</b>	<b>83.3</b>	<b>69.9</b>	<b>79.3</b>	<b>75.8</b>	<b>54.5</b>	<b>63.1</b>	<b>45.2</b>
OSTrack [44]	ECCV'22	71.0	80.4	68.2	83.1	87.8	82.0	69.1	78.7	75.2	-	-	-
AiATrack [11]	ECCV'22	69.6	80.0	63.2	82.7	87.8	80.4	69.0	79.4	73.8	-	-	-
SimTrack [3]	ECCV'22	68.6	78.9	62.4	82.3	86.5	-	69.3	78.5	74.0	-	-	-
RTS [33]	ECCV'22	-	-	-	81.6	86.0	79.4	69.7	76.2	73.7	50.8	55.7	38.9
Unicorn [40]	ECCV'22	-	-	-	83.0	86.4	<b>82.2</b>	68.5	76.6	74.1	-	-	-
MixFormer [5]	CVPR'22	70.7	80.0	67.8	<b>83.1</b>	<b>88.1</b>	81.6	69.2	78.7	74.7	<b>53.7</b>	<b>63.0</b>	<b>43.0</b>
ToMP [28]	CVPR'22	-	-	-	81.2	86.2	78.6	67.6	78.0	72.2	51.6	59.5	38.9
SBT [39]	CVPR'22	69.9	80.4	63.6	-	-	-	65.9	-	70.0	-	-	-
CSWinTT [36]	CVPR'22	69.4	78.9	65.4	81.9	86.7	79.5	66.2	75.2	70.9	-	-	-
STARK [41]	ICCV'21	68.0	77.7	62.3	81.3	86.1	78.1	66.4	76.3	71.2	51.1	59.2	39.1
KeepTrack [29]	ICCV'21	-	-	-	-	-	-	67.1	77.2	70.2	49.4	56.3	37.8
AutoMatch [47]	ICCV'21	65.2	76.6	54.3	76.0	-	72.6	58.3	-	59.9	-	-	-
TransT [4]	CVPR'21	67.1	76.8	60.9	81.4	86.7	80.3	64.9	73.8	69.0	49.0	56.4	37.2
Alpha-Refine [43]	CVPR'21	-	-	-	80.5	85.6	78.3	65.3	73.2	68.0	49.6	55.7	38.2
TMT [38]	CVPR'21	67.1	77.7	58.3	78.4	83.3	73.1	63.9	-	61.4	48.1	55.3	33.8
Ocean [49]	ECCV'20	61.1	72.1	47.3	-	-	-	56.0	65.1	56.6	38.9	43.6	20.5
PrDiMP [7]	CVPR'20	63.4	73.8	54.3	75.8	81.6	70.4	59.8	68.8	60.8	43.3	48.0	28.7
SiamAttn [46]	CVPR'20	-	-	-	75.2	81.7	-	56.0	64.8	-	-	-	-
DiMP [2]	ICCV'19	61.1	71.7	49.2	74.0	80.1	68.7	56.9	65.0	56.7	41.9	45.7	26.0
ATOM [6]	CVPR'19	-	-	-	70.3	77.1	64.8	51.5	57.6	50.5	38.6	41.5	22.2
SiamRPN++ [21]	CVPR'19	51.7	61.6	32.5	73.3	80.0	69.4	49.6	56.9	49.1	39.0	43.5	21.2

Table 2. Comparison with the state-of-the-art trackers on NfS30 and UAV123 in terms of AUC score.

Tracker	SiamRPN++ [21]	ATOM [6]	DiMP [2]	TransT [4]	TMT [38]	STARK [41]	ToMP [28]	OSTrack [44]	GRM
NfS30 [19]	50.2	59.0	62.0	65.7	<b>66.5</b>	65.2	<b>66.9</b>	64.7	65.6
UAV123 [30]	61.3	65.0	65.4	<b>69.1</b>	67.5	<b>69.1</b>	69.0	68.3	<b>70.2</b>

Table 3. Comparison with other performance-oriented variants.

Tracker	LaSOT [9]			TrackingNet [31]		
	AUC	P <sub>Norm</sub>	P	AUC	P <sub>Norm</sub>	P
GRM-L320	<b>71.4</b>	<b>81.2</b>	<b>77.9</b>	<b>84.4</b>	<b>88.9</b>	<b>84.0</b>
OSTrack-384 [44]	71.1	81.1	77.6	83.9	88.5	83.2
Sim-L/14 [3]	70.5	79.7	76.2	83.4	87.4	-
MixFormer-L22k [5]	70.1	79.9	76.3	<b>83.9</b>	<b>88.9</b>	83.1

the network backbone to ViT-L [8] and increase the resolution of the search region to  $320 \times 320$  pixels with  $5^2$  times the target object area. Table 3 illustrates that our variant is highly competitive compared to the performance-oriented variants of most recent works [3, 5, 44].

### 4.3. Ablation Studies

To analyze the effectiveness of the proposed method, we conduct detailed ablation studies on GOT-10k [16]. Our experimental results are summarized in Table 4 and Table 5. **Search Token Division Strategy.** To demonstrate the superiority of our method over previous two-stream and one-stream pipelines, we first devise two counterparts of them using our definition of token categories (Eq. (2), Eq. (3) and Eq. (4)). As shown in Table 4, the two-stream pipeline (variant #1) is equivalent to the situation where all the search tokens are consistently assigned to the category  $E_S$ , and the one-stream pipeline (variant #2) is equivalent to the situation where all the search tokens are consistently assigned to the category  $E_T$ . For a fair comparison, in variant #1, we conduct separate attentions for template tokens and search

Table 4. Ablative experiments of different search token divisions.

#	Token Category			GOT-10k* [16]			Speed fps
	$E_T$	$E_S$	$E_A$	AO	SR <sub>0.5</sub>	SR <sub>0.75</sub>	
1		✓		64.7	73.3	60.2	60
2	✓			71.0	80.3	68.2	51
3	✓	✓		71.5	80.6	68.3	45
4	✓	✓	✓	71.9	81.3	69.0	44
5		✓	✓	<b>73.4</b>	<b>82.9</b>	<b>70.4</b>	45

Table 5. Ablative experiments about where we apply the token division and how we aggregate the target-aware representation.

#	Layer	Aggregation	GOT-10k* [16]		
			AO	SR <sub>0.5</sub>	SR <sub>0.75</sub>
a	None	None	71.0	80.3	68.2
b	1 - 12	Maximum Pooling	71.2	80.6	67.8
c	5 - 12	Maximum Pooling	72.3	81.6	69.3
d	2 - 12	Average Pooling	71.8	81.1	68.9
e	2 - 12	Maximum Pooling	<b>73.4</b>	<b>82.9</b>	<b>70.4</b>

tokens in all encoder layers except for the last one which acts as the correlation module by applying a global attention to all the template and search tokens. Comparing variant #1 with variant #2, we find that the one-stream pipeline performs much better than the two-stream pipeline with the same number of model parameters. It suggests earlier interaction between the template and search region is essential.

We then use the prediction module to classify the search tokens into  $E_T$  and  $E_S$ , which do not interact with each other. Although this division seems to be a possible solution to overcome the limitation of the one-stream pipeline, we find that the improvement of variant #3 over variant #2




Figure 3. Visualization of some representative search token division results in several encoder layers. We choose the  $3^{rd}$ , the  $7^{th}$  and the  $11^{th}$  encoder layer as the examples of the intermediate form, the degenerated two-stream and one-stream forms of our generalized relation modeling. The shaded region denotes the search tokens that are classified as the category  $\mathbf{E}_S$  and do not interact with the template tokens. The ground-truth bounding boxes and the frame indices are shown in the first figure of each line.

is limited since it would block the interaction of some objects with their relevant context in the search region, impeding the feature extraction. In variant #4, the token division module classifies the search tokens into  $\mathbf{E}_T$ ,  $\mathbf{E}_S$  and  $\mathbf{E}_A$ . From the result of variant #4, we can find that the limitation of variant #3 can be alleviated by introducing the third token category  $\mathbf{E}_A$ . However, it still blocks the interaction between the search tokens in  $\mathbf{E}_T$  and those in  $\mathbf{E}_S$ . Finally, by dividing search tokens into  $\mathbf{E}_S$  and  $\mathbf{E}_A$ , the completeness of interaction among the search tokens can be ensured, and the template tokens can still interact with the selected tokens in  $\mathbf{E}_A$  that are suitable for cross-relation modeling. Thus, the overall performance of variant #5 is the best among all variants in Table 4.

**Encoder Layers with Token Division.** To investigate which encoder layers are suitable to conduct adaptive token division, we implement three variants of our method by inserting the prediction module we proposed in Sec. 3.3 to all encoder layers (variant #b), the last eight encoder layers (variant #c), and all encoder layers except for the first one (variant #e). From Table 5, we can see that both variant #b and variant #c are inferior to variant #e. The result of variant #b suggests that it is inappropriate to perform token division by the prediction module in the first encoder layer where the raw embedding projections have not been refined

by any feature aggregation. The result of variant #c suggests that our token division method can perform favorably in other earlier encoder layers. Therefore, we conduct adaptive token division in all encoder layers except for the first layer as the default setting in our experiments.

**Target-Aware Representation Aggregation.** We also explore the aggregation method to generate target-aware representation, which is used to predict the category of search tokens (Eq. (6)). From the last two rows of Table 5, we can see that aggregation via maximum pooling (variant #e) is better than average pooling (variant #d). The underlying reason might be that the maximum pooling operation is less sensitive to noise, which may be introduced by the surrounding background inside the template image.

**Speedup by Attention Masking.** As mentioned in Sec. 3.3, separate attention operations will become a bottleneck for the running speed. For a quantitative comparison, we evaluate the inference speed of the two implementations. The speed with separate attention operations is 33 fps, whilst the proposed attention masking strategy accelerates it to 45 fps.

#### 4.4. Visualizations

To investigate how our generalized relation modeling method works in each encoder layer, we visualize some representative search token division results in Fig. 3. From the visualization results, we observe that: (i) in the  $3^{rd}$  encoder layer, our prediction module classifies the search tokens relevant to the target as  $\mathbf{E}_A$  to perform cross-relation modeling with the template tokens; (ii) in the  $7^{th}$  encoder layer, none of the search tokens are selected as  $\mathbf{E}_A$  and the relation modeling in this layer degenerates to a two-stream form; (iii) in the  $11^{th}$  encoder layer, all of the search tokens are selected as  $\mathbf{E}_A$  and it degenerates to a one-stream form. In fact, we find that the earlier layers tend to perform as the two-stream form or the intermediate form in which a portion of search tokens are selected to interact with the template tokens. Besides, the relation modeling degenerates to the one-stream form in the last five encoder layers. This suggests that highly discriminative representations have been extracted in the deeper layers, and the constraints on global interaction are not necessary for these layers. More thorough discussions can be found in the appendices.

## 5. Conclusion

In this paper, we present a generalized relation modeling method, which inherits the strengths of both the two-stream and one-stream Transformer tracking pipelines whilst being more flexible. An end-to-end optimized token division module with attention masking strategy is introduced to perform adaptive token division. Extensive experiments and analyses demonstrate the efficacy of the proposed method. We hope this work would shed some new light on promoting the relation modeling for Transformer trackers.



## References

- [1] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *European conference on computer vision*, pages 850–865. Springer, 2016. 1, 2
- [2] 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. 2, 7
- [3] 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. *arXiv preprint arXiv:2203.05328*, 2022. 1, 3, 7
- [4] 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. 1, 2, 7
- [5] 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. 1, 3, 6, 7
- [6] 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. 2, 7
- [7] Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7183–7192, 2020. 7
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 6, 7
- [9] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. Lasot: A high-quality benchmark for large-scale single object tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5374–5383, 2019. 6, 7
- [10] Zhihong Fu, Zehua Fu, Qingjie Liu, Wenrui Cai, and Yunhong Wang. Sparset: Visual tracking with sparse transformers. *arXiv preprint arXiv:2205.03776*, 2022. 3
- [11] Shenyuan Gao, Chunluan Zhou, Chao Ma, Xinggang Wang, and Junsong Yuan. Aiatrack: Attention in attention for transformer visual tracking. In *European Conference on Computer Vision*, pages 146–164. Springer, 2022. 1, 2, 7
- [12] Mingzhe Guo, Zhipeng Zhang, Heng Fan, Liping Jing, Yilin Lyu, Bing Li, and Weiming Hu. Learning target-aware representation for visual tracking via informative interactions. *arXiv preprint arXiv:2201.02526*, 2022. 2
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 3, 6
- [14] 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. 2
- [15] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 6
- [16] 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, 7
- [17] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. 2, 5
- [18] Sajid Javed, Martin Danelljan, Fahad Shahbaz Khan, Muhammad Haris Khan, Michael Felsberg, and Jiri Matas. Visual object tracking with discriminative filters and siamese networks: a survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1
- [19] Hamed Kiani Galoogahi, Ashton Fagg, Chen Huang, Deva Ramanan, and Simon Lucey. Need for speed: A benchmark for higher frame rate object tracking. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1125–1134, 2017. 6, 7
- [20] Zhenglun Kong, Peiyan Dong, Xiaolong Ma, Xin Meng, Wei Niu, Mengshu Sun, Bin Ren, Minghai Qin, Hao Tang, and Yanzhi Wang. Spvit: Enabling faster vision transformers via soft token pruning. *arXiv preprint arXiv:2112.13890*, 2021. 3
- [21] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4282–4291, 2019. 1, 2, 7
- [22] 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. 1, 2
- [23] Liting Lin, Heng Fan, Yong Xu, and Haibin Ling. Swintrack: A simple and strong baseline for transformer tracking. *arXiv preprint arXiv:2112.00995*, 2021. 2
- [24] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 6
- [25] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 3

- [27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [28] Christoph Mayer, Martin Danelljan, Goutam Bhat, Matthieu Paul, Danda Pani Paudel, Fisher Yu, and Luc Van Gool. Transforming model prediction for tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8731–8740, 2022. 2, 7
- [29] Christoph Mayer, Martin Danelljan, Danda Pani Paudel, and Luc Van Gool. Learning target candidate association to keep track of what not to track. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13444–13454, 2021. 7
- [30] Matthias Mueller, Neil Smith, and Bernard Ghanem. A benchmark and simulator for uav tracking. In *European conference on computer vision*, pages 445–461. Springer, 2016. 6, 7
- [31] 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, 7
- [32] Mubashir Noman, Wafa Al Ghallabi, Daniya Najiha, Christoph Mayer, Akshay Dudhane, Martin Danelljan, Hisham Cholakkal, Salman Khan, Luc Van Gool, and Fahad Shahbaz Khan. Avist: A benchmark for visual object tracking in adverse visibility. *arXiv preprint arXiv:2208.06888*, 2022. 6, 7
- [33] Matthieu Paul, Martin Danelljan, Christoph Mayer, and Luc Van Gool. Robust visual tracking by segmentation. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, pages 571–588. Springer, 2022. 7
- [34] 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. 3, 5
- [35] Hamid Rezaatofghi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666, 2019. 6
- [36] Zikai Song, Junqing Yu, Yi-Ping Phoebe Chen, and Wei Yang. Transformer tracking with cyclic shifting window attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8791–8800, 2022. 7
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [38] Ning Wang, Wengang Zhou, Jie Wang, and Houqiang Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1571–1580, 2021. 1, 2, 7
- [39] Fei Xie, Chunyu Wang, Guangting Wang, Yue Cao, Wankou Yang, and Wenjun Zeng. Correlation-aware deep tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8751–8760, 2022. 1, 2, 7
- [40] Bin Yan, Yi Jiang, Peize Sun, Dong Wang, Zehuan Yuan, Ping Luo, and Huchuan Lu. Towards grand unification of object tracking. *arXiv preprint arXiv:2207.07078*, 2022. 2, 7
- [41] 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. 2, 7
- [42] 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. 3
- [43] Bin Yan, Xinyu Zhang, Dong Wang, Huchuan Lu, and Xiaoyun Yang. Alpha-refine: Boosting tracking performance by precise bounding box estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5289–5298, 2021. 7
- [44] 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, 6, 7
- [45] Tan Yu, Gangming Zhao, Ping Li, and Yizhou Yu. Boat: Bilateral local attention vision transformer. *arXiv preprint arXiv:2201.13027*, 2022. 3
- [46] Yuechen Yu, Yilei Xiong, Weilin Huang, and Matthew R Scott. Deformable siamese attention networks for visual object tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6728–6737, 2020. 2, 7
- [47] Zhipeng Zhang, Yihao Liu, Xiao Wang, Bing Li, and Weiming Hu. Learn to match: Automatic matching network design for visual tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13339–13348, 2021. 3, 7
- [48] Zhipeng Zhang and Houwen Peng. Deeper and wider siamese networks for real-time visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4591–4600, 2019. 2
- [49] Zhipeng Zhang, Houwen Peng, Jianlong Fu, Bing Li, and Weiming Hu. Ocean: Object-aware anchor-free tracking. In *European Conference on Computer Vision*, pages 771–787. Springer, 2020. 1, 2, 7