

Transformer Meets Tracker: Exploiting Temporal Context for Robust Visual Tracking

Ning Wang¹ Wengang Zhou^{1,2} Jie Wang^{1,2} Houqiang Li^{1,2}

¹CAS Key Laboratory of GIPAS, EEIS Department, University of Science and Technology of China (USTC)

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

wn6149@mail.ustc.edu.cn, {zhwg, jiewangx, lihq}@ustc.edu.cn

Abstract

In video object tracking, there exist rich temporal contexts among successive frames, which have been largely overlooked in existing trackers. In this work, we bridge the individual video frames and explore the temporal contexts across them via a transformer architecture for robust object tracking. Different from classic usage of the transformer in natural language processing tasks, we separate its encoder and decoder into two parallel branches and carefully design them within the Siamese-like tracking pipelines. The transformer encoder promotes the target templates via attention-based feature reinforcement, which benefits the high-quality tracking model generation. The transformer decoder propagates the tracking cues from previous templates to the current frame, which facilitates the object searching process. Our transformer-assisted tracking framework is neat and trained in an end-to-end manner. With the proposed transformer, a simple Siamese matching approach is able to outperform the current top-performing trackers. By combining our transformer with the recent discriminative tracking pipeline, our method sets several new state-of-the-art records on prevalent tracking benchmarks.

1. Introduction

Visual object tracking is a basic task in computer vision. Despite the recent progress, it remains a challenging task due to factors such as occlusion, deformation, and appearance changes. With the temporal error accumulation, these challenges are further amplified in the online process.

It is well recognized that the rich temporal information in the video flow is of vital importance for visual tracking. However, most tracking paradigms [28, 27, 48] handle this task by per-frame object detection, where the tempo-




Figure 1. An overview of our transformer-assisted tracking framework. The transformer encoder and decoder are assigned to two parallel branches in a Siamese-like tracking pipeline. Thanks to the encoder-decoder structure, isolated frames are tightly bridged to convey rich temporal information in the video flow.

ral relationships among successive frames have been largely overlooked. Take the popular Siamese tracker as an example, only the initial target is considered for template matching [1, 44, 19, 28]. The merely used temporal information is the motion prior (e.g., cosine window) by assuming the target moves smoothly, which is widely adopted in visual trackers. In other tracking frameworks with update mechanisms [20, 39, 8, 58, 60, 3], previous prediction results are collected to incrementally update the tracking model. Despite the historical frames considered in the above approaches, the video frames are still considered as independent counterparts without mutual reasoning. In real-world videos, some frames inevitably contain noisy contents such as occluded or blurred objects. These imperfect frames will hurt the model update when serving as the templates and will challenge the tracking process when performing as the search frames. Therefore, it is a non-trivial issue to convey rich information across temporal frames to mutually reinforce them. We argue that the video frames should not be treated in isolation and the performance potential is largely restricted due to the overlook of frame-wise relationship.

To bridge the isolated video frames and convey the rich temporal cues across them, in this work, we introduce the transformer architecture [46] to the visual tracking community. Different from the traditional usage of the transformer

*Corresponding Author: Wengang Zhou and Houqiang Li.

†Source code, pretrained model, and raw tracking results are available at <https://github.com/594422814/TransformerTrack>.

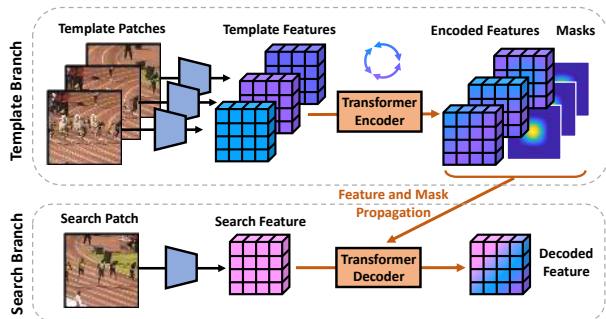


Figure 2. Top: the transformer encoder receives multiple template features to mutually aggregate representations. Bottom: the transformer decoder propagates the template features and their assigned masks to the search patch feature for representation enhancement.

in language modeling and machine translation [46, 12], we leverage it to handle the *context propagation in the temporal domain*. By carefully modifying the classic transformer architecture, we show that its transformation characteristic naturally fits the tracking scenario. Its core component, *i.e.*, attention mechanism [46, 55], is ready to establish the pixel-wise correspondence across frames and freely convey various signals in the temporal domain.

Generally, most tracking methods [1, 45, 28, 42, 7, 3] can be formulated into a Siamese-like framework, where the top branch learns a tracking model using template features, and the bottom branch classifies the current search patch. As shown in Figure 1, we separate the transformer encoder and decoder into two branches within such a general Siamese-like structure. In the top branch, a set of template patches are fed to the transformer encoder to generate high-quality encoded features. In the bottom branch, the search feature as well as the previous template contents are fed to the transformer decoder, where the search patch retrieves and aggregates informative target cues (*e.g.*, spatial masks and target features) from history templates to reinforce itself.

The proposed transformer facilitates visual tracking via:

- **Transformer Encoder.** It enables individual template features to mutually reinforce to acquire more compact target representations, as shown in Figure 2. These encoded high-quality features further benefit the tracking model generation.
- **Transformer Decoder.** It conveys valuable temporal information across frames. As shown in Figure 2, our decoder simultaneously transfers features and spatial masks. Propagating the features from previous frames to the current patch smooths the appearance changes and remedies the context noises while transforming the spatial attentions highlights the potential object location. These manifold target representations and spatial cues make the object search much easier.

Finally, we track the target in the decoded search patch. To verify the generalization of our designed transformer, we

integrate it into two popular tracking frameworks including a Siamese formulation [1] and a discriminative correlation filter (DCF) based tracking paradigm [3]. With our designed transformer, a simple Siamese matching pipeline is able to outperform the current top-performing trackers. By combining with the recent discriminative approach [3], our transformer-assisted tracker shows outstanding results on seven prevalent tracking benchmarks including LaSOT [13], TrackingNet [38], GOT-10k [23], UAV123 [36], NfS [24], OTB-2015 [56], and VOT2018 [26] and sets several new state-of-the-art records.

In summary, we make three-fold contributions:

- We present a neat and novel transformer-assisted tracking framework. To our best knowledge, this is the first attempt to involve the transformer in visual tracking.
- We simultaneously consider the feature and attention transformations to better explore the potential of the transformer. We also modify the classic transformer to make it better suit the tracking task.
- To verify the generalization, we integrate our designed transformer into two popular tracking pipelines. Our trackers exhibit encouraging results on 7 benchmarks.

2. Related Work

Visual Tracking. Given the initial target in the first frame, visual tracking aims to localize it in successive frames. In recent years, the Siamese network has gained significant popularity, which deals with the tracking task by template matching [1, 44, 19]. By introducing the region proposal network (RPN), Siamese trackers obtain superior efficiency and more accurate target scale estimation [28, 63]. The recent improvements upon Siamese trackers include attention mechanism [54], reinforcement learning [22, 51], target-aware model fine-tuning [30], unsupervised training [50, 52], sophisticated backbone networks [27, 61], cascaded frameworks [14, 49], and model update mechanisms [16, 17, 58, 60].

Discriminative correlation filter (DCF) tackles the visual tracking by solving the ridge regression in Fourier domain, which exhibits attractive efficiency [20, 35, 34, 15, 37, 53, 11, 8]. The recent advances show that the ridge regression can be solved in the deep learning frameworks [42, 32, 7, 3], which avoids the boundary effect in classic DCF trackers. These methods learn a discriminative CNN kernel to convolve with the search area for response generation. In recent works, the residual terms [42] and shrinkage loss [32] are incorporated into the deep DCF formulation. To accelerate the kernel learning process, ATOM [7] exploits the conjugate gradient algorithm. The recent DiMP tracker [3] enhances the discriminative capability of the learned CNN kernel in an end-to-end manner, which is further promoted by the probabilistic regression framework [9].

Despite the impressive performance, most existing meth-

ods [39, 1, 28, 7, 3, 33, 48] generally regard the tracking task as the per-frame object detection problem, failing to adequately exploit the temporal characteristic of the tracking task. Some previous works explore the temporal information using graph neural network [16], spatial-temporal regularization [29], optical flow [64], *etc.* Differently, we leverage the transformer to model the frame-wise relationship and propagate the temporal cues, which is neat and ready to integrate with the modern deep trackers.

Transformer. Transformer is first proposed in [46] as a new paradigm for machine translation. The basic block in a transformer is the attention module, which aggregates information from the entire input sequence. Due to the parallel computations and unique memory mechanism, transformer architecture is more competitive than RNNs in processing long sequences and has gained increasing popularity in many natural language processing (NLP) tasks [12, 41, 43]. Similarly, non-local neural network [55] also introduces a self-attention block to acquire global representations, which has been adopted in many vision tasks including visual object tracking [59]. Nevertheless, how to take advantage of the compact transformer encoder-decoder structure for visual tracking has been rarely studied.

Recently, transformer architecture has been introduced to computer vision such as image generation [40]. Transformer based object detection approach is proposed in [5], which views the object detection task as a direct set prediction problem. However, the above techniques leverage the transformer in the image-level tasks. In this paper, we show that the transformer structure serves as a good fit for video-related scenarios by transferring temporal information across frames. To bridge the domain gap between visual tracking and NLP tasks, we carefully modify the classic transformer to better suit the tracking scenario.

3. Revisiting Tracking Frameworks

Before elaborating our transformer for object tracking, we briefly review the recent popular tracking approaches for the sake of completeness. As shown in Figure 3, the mainstream tracking methods such as Siamese network [1] or discriminative correlation filter (DCF) [45, 7, 3] can be formulated into the Siamese-like pipeline, where the top branch learns the tracking model using templates and the bottom branch focuses on the target localization.

Siamese matching architecture [1] takes an exemplar patch z and a search patch x as inputs, where z represents the target object while x is a large searching area in subsequent video frames. Both of them are fed to the weight-sharing CNN network $\Psi(\cdot)$. Their output feature maps are cross-correlated as follows to generate the response map:

$$r(z, x) = \Psi(z) * \Psi(x) + b \cdot \mathbb{1}, \quad (1)$$

where $*$ is the cross-correlation and $b \cdot \mathbb{1}$ denotes a bias

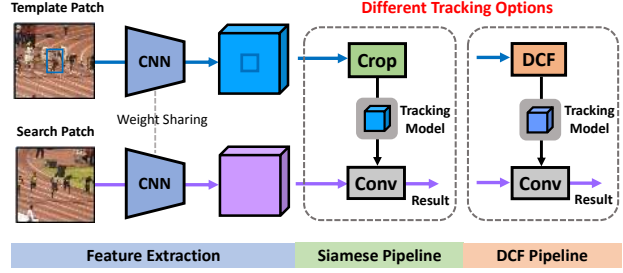


Figure 3. The simplified pipelines of Siamese [1] and DCF [7, 3] based trackers. These tracking approaches can be formulated into a Siamese-like pipeline, where the top branch is responsible for the model generation and the bottom branch localizes the target.

term. Siamese trackers rely on the target model, *i.e.*, convolutional kernel $\Psi(z)$, for template matching.

As another popular framework, deep learning based DCF method optimizes the tracking model f under a ridge regression formulation [42, 7, 3] as follows:

$$\min_f \|f * \Psi(z^*) - y\|_2^2 + \lambda \|f\|_2^2, \quad (2)$$

where y is the Gaussian-shaped ground-truth label of template patch z^* , and λ controls the regularization term to avoid overfitting. Note that z^* is much larger than the exemplar patch z in Siamese trackers. Therefore, DCF formulation simultaneously considers the target matching and background discrimination. After obtaining the tracking model f , the response is generated via $r = f * \Psi(x)$.

The traditional DCF methods [20, 10] solve ridge regression using circularly generated samples via the closed-form solution in the Fourier domain. In contrast, the recent deep learning based DCF methods solve Eq. 2 using stochastic gradient descent [42, 32] or conjugate gradient approach [7] to avoid the boundary effect. The recent DiMP [3] optimizes the above ridge regression via a meta-learner in an end-to-end manner, showing state-of-the-art performance.

4. Transformer for Visual Tracking

As discussed in Section 3, mainstream tracking methods can be formulated into a Siamese-like pipeline. We aim to improve such a general tracking framework by frame-wise relationship modeling and temporal context propagation, without modifying their original tracking manners such as template matching.

4.1. Transformer Overview

An overview of our transformer is shown in Figure 4. Similar to the classic transformer architecture [46], the encoder leverages self-attention block to mutually reinforce multiple template features. In the decoding process, cross-attention block bridges template and search branches to propagate temporal contexts (*e.g.*, feature and attention).




Figure 4. An overview of the proposed transformer architecture.

To suit the visual tracking task, we modify the classic transformer in the following aspects: (1) *Encoder-decoder Separation*. Instead of cascading the encoder and decoder in NLP tasks [46, 12], as shown in Figure 1, we separate the encoder and decoder into two branches to fit the Siamese-like tracking methods. (2) *Block Weight-sharing*. The self-attention blocks in the encoder and decoder (yellow boxes in Figure 4) share weights, which transform the template and search embeddings in the same feature space to facilitate the further cross-attention computation. (3) *Instance Normalization*. In NLP tasks [46], the word embeddings are individually normalized using the layer normalization. Since our transformer receives image feature embeddings, we jointly normalize these embeddings in the instance (image patch) level to retain the valuable image amplitude information. (4) *Slimming Design*. Efficiency is crucial for visual tracking scenarios. To achieve a good balance of speed and performance, we slim the classic transformer by omitting the fully-connected feed-forward layers and maintaining the lightweight single-head attention.

4.2. Transformer Encoder

The basic block in a classic transformer is the attention mechanism, which receives the query $\mathbf{Q} \in \mathbb{R}^{N_q \times C}$, key $\mathbf{K} \in \mathbb{R}^{N_k \times C}$, and value $\mathbf{V} \in \mathbb{R}^{N_k \times C}$ as the inputs. In our approach, following [46], we also adopt the dot-product to compute the similarity matrix $\mathbf{A}_{\mathbf{K} \rightarrow \mathbf{Q}} \in \mathbb{R}^{N_q \times N_k}$ between the query and key as follows:

$$\mathbf{A}_{\mathbf{K} \rightarrow \mathbf{Q}} = \text{Atten}(\mathbf{Q}, \mathbf{K}) = \text{Softmax}_{\text{col}}(\bar{\mathbf{Q}}\bar{\mathbf{K}}^T/\tau), \quad (3)$$

where $\bar{\mathbf{Q}}$ and $\bar{\mathbf{K}}$ are ℓ_2 -normalized features of \mathbf{Q} and \mathbf{K} across the channel dimension, and τ is a temperature parameter controlling the Softmax distribution, which is inspired by the model distillation [21] and contrastive learning [6] techniques. With the propagation matrix $\mathbf{A}_{\mathbf{K} \rightarrow \mathbf{Q}}$ from key to query, we can transform the value via $\mathbf{A}_{\mathbf{K} \rightarrow \mathbf{Q}}\mathbf{V} \in \mathbb{R}^{N_q \times C}$.

In our framework, the transformer encoder receives a set of template features $\mathbf{T}_i \in \mathbb{R}^{C \times H \times W}$ with a spatial size of $H \times W$ and dimensionality C , which are further concatenated to form the template feature ensemble $\mathbf{T} = \text{Concat}(\mathbf{T}_1, \dots, \mathbf{T}_n) \in \mathbb{R}^{n \times C \times H \times W}$. To facilitate the attention computation, we reshape \mathbf{T} to $\mathbf{T}' \in \mathbb{R}^{N_T \times C}$, where $N_T = n \times H \times W$. As shown in Figure 4, the main operation in the transformer encoder is self-attention, which aims to mutually reinforce the features from multiple templates. To this end, we first compute the self-attention map $\mathbf{A}_{\mathbf{T} \rightarrow \mathbf{T}} = \text{Atten}(\varphi(\mathbf{T}'), \varphi(\mathbf{T}')) \in \mathbb{R}^{N_T \times N_T}$, where $\varphi(\cdot)$ is a 1×1 linear transformation that reduces the embedding channel from C to $C/4$.

Based on the self-similarity matrix $\mathbf{A}_{\mathbf{T} \rightarrow \mathbf{T}}$, we transform the template feature through $\mathbf{A}_{\mathbf{T} \rightarrow \mathbf{T}}\mathbf{T}'$, which is added to the original feature \mathbf{T}' as a residual term as follows:

$$\hat{\mathbf{T}} = \text{Ins. Norm}(\mathbf{A}_{\mathbf{T} \rightarrow \mathbf{T}}\mathbf{T}' + \mathbf{T}'), \quad (4)$$

where $\hat{\mathbf{T}} \in \mathbb{R}^{N_T \times C}$ is the encoded template feature and $\text{Ins. Norm}(\cdot)$ denotes the instance normalization that jointly ℓ_2 -normalizes all the embeddings from an image patch, *i.e.*, feature map level ($\mathbf{T}_i \in \mathbb{R}^{C \times H \times W}$) normalization.

Thanks to the self-attention, multiple temporally diverse template features aggregate each other to generate high-quality $\hat{\mathbf{T}}$, which is further fed to the decoder block to reinforce the search patch feature. Besides, this encoded template representation $\hat{\mathbf{T}}$ is also reshaped back to $\mathbf{T}^{\text{encoded}} \in \mathbb{R}^{n \times C \times H \times W}$ for tracking model generation, *e.g.*, the DCF model in Section 4.4.

4.3. Transformer Decoder

Transformer decoder takes the search patch feature $\mathbf{S} \in \mathbb{R}^{C \times H \times W}$ as its input. Similar to the encoder, we first reshape this feature to $\mathbf{S}' \in \mathbb{R}^{N_S \times C}$, where $N_S = H \times W$. Then, \mathbf{S}' is fed to the self-attention block as follows:

$$\hat{\mathbf{S}} = \text{Ins. Norm}(\mathbf{A}_{\mathbf{S} \rightarrow \mathbf{S}}\mathbf{S}' + \mathbf{S}'), \quad (5)$$

where $\mathbf{A}_{\mathbf{S} \rightarrow \mathbf{S}} = \text{Atten}(\varphi(\mathbf{S}'), \varphi(\mathbf{S}')) \in \mathbb{R}^{N_S \times N_S}$ is the self-attention matrix of the search feature.

Mask Transformation. Based on the search feature $\hat{\mathbf{S}}$ in Eq. 5 and aforementioned encoded template feature $\hat{\mathbf{T}}$ in Eq. 4, we compute the cross-attention matrix between them via $\mathbf{A}_{\mathbf{T} \rightarrow \mathbf{S}} = \text{Atten}(\phi(\hat{\mathbf{S}}), \phi(\hat{\mathbf{T}})) \in \mathbb{R}^{N_S \times N_T}$, where

$\phi(\cdot)$ is a 1×1 linear transformation block similar to $\varphi(\cdot)$. This cross-attention map $\mathbf{A}_{T \rightarrow S}$ establishes the pixel-to-pixel correspondence between frames, which supports the temporal context propagation.

In visual tracking, we are aware of the target positions in the templates. To propagate the temporal motion priors, we construct the Gaussian-shaped masks of the template features through $\mathbf{m}(y) = \exp\left(-\frac{\|y-c\|^2}{2\sigma^2}\right)$, where c is the ground-truth target position. Similar to the feature ensemble \mathbf{T} , we also concatenate these masks $\mathbf{m}_i \in \mathbb{R}^{H \times W}$ to form the mask ensemble $\mathbf{M} = \text{Concat}(\mathbf{m}_1, \dots, \mathbf{m}_n) \in \mathbb{R}^{n \times H \times W}$, which is further flattened into $\mathbf{M}' \in \mathbb{R}^{N_T \times 1}$. Based on the cross attention map $\mathbf{A}_{T \rightarrow S}$, we can easily propagate previous masks to the search patch via $\mathbf{A}_{T \rightarrow S} \mathbf{M}' \in \mathbb{R}^{N_S \times 1}$. The transformed mask is qualified to serve as the attention weight for the search feature $\hat{\mathbf{S}}$ as follows:

$$\hat{\mathbf{S}}_{\text{mask}} = \text{Ins. Norm} \left(\mathbf{A}_{T \rightarrow S} \mathbf{M}' \otimes \hat{\mathbf{S}} \right), \quad (6)$$

where \otimes is the broadcasting element-wise multiplication. By virtue of the spatial attention, the reinforced search feature $\hat{\mathbf{S}}_{\text{mask}}$ better highlights the potential target area.

Feature Transformation. Except for the spatial attention, it is also feasible to propagate the context information from template feature $\hat{\mathbf{T}}$ to the search feature $\hat{\mathbf{S}}$. It is beneficial to convey target representations while the background scenes tend to change drastically in a video, which is unreasonable to temporally propagate. As a consequence, before feature transformation, we first mask the template feature through $\hat{\mathbf{T}} \otimes \mathbf{M}'$ to suppress the background area. Then, with the cross-attention matrix $\mathbf{A}_{T \rightarrow S}$, the transformed feature can be computed via $\mathbf{A}_{T \rightarrow S}(\hat{\mathbf{T}} \otimes \mathbf{M}') \in \mathbb{R}^{N_S \times C}$, which is added to $\hat{\mathbf{S}}$ as a residual term:

$$\hat{\mathbf{S}}_{\text{feat}} = \text{Ins. Norm} \left(\mathbf{A}_{T \rightarrow S}(\hat{\mathbf{T}} \otimes \mathbf{M}') + \hat{\mathbf{S}} \right). \quad (7)$$

Compared with original $\hat{\mathbf{S}}$, feature-level enhanced $\hat{\mathbf{S}}_{\text{feat}}$ aggregates temporally diverse target representations from a series of template features $\hat{\mathbf{T}}$ to promote itself.

Finally, we equally combine the aforementioned spatially masked feature $\hat{\mathbf{S}}_{\text{mask}}$ and feature-level enhanced feature $\hat{\mathbf{S}}_{\text{feat}}$, and further normalize them as follows:

$$\hat{\mathbf{S}}_{\text{final}} = \text{Ins. Norm} \left(\hat{\mathbf{S}}_{\text{feat}} + \hat{\mathbf{S}}_{\text{mask}} \right). \quad (8)$$

The final output feature $\hat{\mathbf{S}}_{\text{final}} \in \mathbb{R}^{N_S \times C}$ is reshaped back to the original size for visual tracking. We denote the reshaped version of $\hat{\mathbf{S}}_{\text{final}}$ as $\mathbf{S}_{\text{decoded}} \in \mathbb{R}^{C \times H \times W}$.

4.4. Tracking with Transformer-enhanced Features

Transformer structure facilitates the tracking process by generating high-quality template feature $\mathbf{T}_{\text{encoded}}$ and search feature $\mathbf{S}_{\text{decoded}}$. We learn the tracking model using $\mathbf{T}_{\text{encoded}}$ following two popular paradigms:

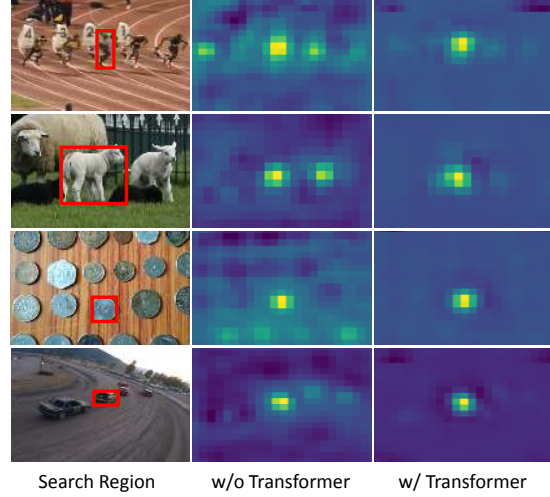


Figure 5. Tracking response maps of the DiMP baseline [3] without (second column) and with (third column) our designed transformer architecture. With the proposed transformer, the confidences of the distracting objects are effectively suppressed.

- **Siamese Pipeline.** In this setting, we simply crop the target feature in $\mathbf{T}_{\text{encoded}}$ as the template CNN kernel to convolve with $\mathbf{S}_{\text{decoded}}$ for response generation, which is identical to the cross-correlation in SiamFC [1].
- **DCF Pipeline.** Following the end-to-end DCF optimization in DiMP approach [3], we generate a discriminative CNN kernel using $\mathbf{T}_{\text{encoded}}$ to convolve with $\mathbf{S}_{\text{decoded}}$ for response generation.

After obtaining the tracking response, we utilize the classification loss proposed in DiMP [3] to jointly train the backbone network, our transformer, and the tracking model in an end-to-end manner. Please refer to [3] for more details.

In the online tracking process, to better exploit the temporal cues and adapt to the target appearance changes, we dynamically update the template ensemble \mathbf{T} . To be specific, we drop the oldest template in \mathbf{T} and add the current collected template feature to \mathbf{T} every 5 frames. The feature ensemble maintains a maximal size of 20 templates. Once the template ensemble \mathbf{T} is updated, we compute the new encoded feature $\mathbf{T}_{\text{encoded}}$ via our transformer encoder. While the transformer encoder is sparsely utilized (*i.e.*, every 5 frames), the transformer decoder is leveraged in each frame, which generates per-frame $\mathbf{S}_{\text{decoded}}$ by propagating the representations and attention cues from previous templates to the current search patch.

It is widely recognized that DCF formulation in DiMP [3] is superior to the simple cross-correlation in Siamese trackers [1, 27]. Nevertheless, in the experiments, we show that with the help of our transformer architecture, a classic Siamese pipeline is able to perform against the recent DiMP. Meanwhile, with our transformer, the DiMP tracker acquires further performance improvements. As shown in

Figure 5, even though the strong baseline DiMP [3] already shows impressive distractor discrimination capability, our designed transformer further assists it to restrain the background confidence for robust tracking.

5. Experiments

5.1. Implementation Details

Based on the Siamese matching and DiMP based tracking frameworks, in the following experiments, we denote our Transformer-assisted trackers as **TrSiam** and **TrDiMP**, respectively. In these two versions, the backbone model is ResNet-50 [18] for feature extraction. Before the encoder and decoder, we additionally add one convolutional layer (3×3 Conv + BN) to reduce the backbone feature channel from 1024 to 512. The input template and search patches are 6 times of the target size and further resized to 352×352 . The temperature τ in Eq. 3 is set to $1/30$. The parameter sigma σ in the feature mask is set to 0.1. Similar to the previous works [7, 3, 9, 2], we utilize the training splits of LaSOT [13], TrackingNet [38], GOT-10k [23], and COCO [31] for offline training. The proposed transformer network is jointly trained with the original tracking parts (*e.g.*, tracking optimization model [3] and IoUNet [9]) in an end-to-end manner. Our framework is trained for 50 epochs with 1500 iterations per epoch and 36 image pairs per batch. The ADAM optimizer [25] is employed with an initial learning rate of 0.01 and a decay factor of 0.2 for every 15 epochs.

In the online tracking stage, the main difference between TrSiam and TrDiMP lies in the tracking model generation manner. After predicting the response map for target localization, they all adopt the recent probabilistic IoUNet [9] for target scale estimation. Our trackers are implemented in Python using PyTorch. TrSiam and TrDiMP operate about 35 and 26 frames per second (FPS) on a single Nvidia GTX 1080Ti GPU, respectively.

5.2. Ablation Study

To verify the effectiveness of our designed transformer structure, we choose the GOT-10k test set [23] with 180 videos to validate our TrSiam and TrDiMP methods¹. GOT-10k hides the ground-truth labels of the test set to avoid the overly hyper-parameter fine-tuning. It is worth mentioning that there is no overlap in object classes between the train and test sets of GOT-10k, which also verifies the generalization of our trackers to unseen object classes.

In Table 1, based on the Siamese and DiMP baselines, we validate each component in our transformer:

Transformer Encoder. First, without any decoder block, we merely utilize encoder to promote the feature fusion of

¹With the probabilistic IoUNet [9] and a larger search area, our baseline performance is better than the standard DiMP [3]. Note that all the experiments (Figure 6 and Table 1) are based on the same baseline for fair.




Figure 6. Training loss plots of the Siamese pipeline (left) and DCF pipeline (right). By combining both feature and mask transformations, our approach significantly reduces the training losses.

Table 1. Ablative experiments of our transformer for the Siamese and DiMP pipelines *i.e.*, TrSiam and TrDiMP trackers. The performance is evaluated on the GOT-10k test set [23] in terms of average overlap (AO).

| Different Tracking Variations | Siamese (AO) | DiMP (AO) |
|--------------------------------------------|-----------------------------|-----------------------------|
| Baseline Performance | 62.0 | 66.7 |
| Only Encoder (w/o Any Decoder) | 63.8 _{1.8%↑} | 67.3 _{0.6%↑} |
| Encoder + Decoder (Only Feature Transf.) | 66.3 _{4.3%↑} | 68.1 _{1.4%↑} |
| Encoder + Decoder (Only Mask Transf.) | 67.1 _{5.1%↑} | 67.8 _{1.1%↑} |
| Encoder + Decoder (Feature & Mask Transf.) | 67.3_{5.3%↑} | 68.8_{2.1%↑} |

Table 2. Ablative study of our transformer architecture. The baseline tracker is TrSiam. The evaluation metric is average overlap (AO) score on the GOT-10k test set.

| | Baseline | Weight-sharing | | Feed-forward | | Head Number | | |
|-------------|----------|----------------|-------------|--------------|------|-------------|------|-------------|
| | | w/o | w/ | w/o | w/ | 1 | 2 | 4 |
| AO (%) | 62.0 | 63.4 | 67.3 | 67.3 | 67.0 | 67.3 | 67.2 | 67.6 |
| Speed (FPS) | 40 | 35 | 35 | 35 | 22 | 35 | 31 | 25 |

multiple templates, which slightly improves two baselines.

Transformer Decoder. Our decoder consists of feature and mask transformations, and we independently verify them: (1) *Feature Propagation.* With the feature transformation, as shown in Table 1, the Siamese pipeline obtains a notable performance gain of 4.3% in AO and the strong DiMP baseline still acquires an improvement of 1.4% in AO on the GOT-10k test set. From the training perspective, we can observe that this block effectively reduces the losses of two baselines as shown in Figure 6.

(2) *Mask Propagation.* This mechanism propagates temporally collected spatial attentions to highlight the target area. Similar to the feature transformation, our mask transformation also steadily improves the tracking performance (Table 1) and consistently reduces the training errors of both two pipelines (Figure 6).

Complete Transformer. With the complete transformer, as shown in Table 1, the Siamese and DiMP baselines obtain notable performance gains of 5.3% and 2.1% in AO, respectively. The transformer also significantly reduces their training losses (Figure 6). It is worth mentioning that DiMP already achieves outstanding results while our approach consistently improves such a strong baseline. With our transformer, the performance gap between Siamese and DiMP

Table 3. Comparison with state-of-the-art trackers on the TrackingNet test set [38] in terms of precision (Prec.), normalized precision (N. Prec.), and success (AUC score). Our TrDiMP and TrSiam exhibit promising results.

| | SiamFC | MDNet | SPM | C-RPN | SiamRPN++ | ATOM | DiMP-50 | SiamFC++ | D3S | Retain-MAML | PrDiMP-50 | DCFST | KYS | Siam-RCNN | TrSiam | TrDiMP |
|--------------|--------|-------|------|-------|-----------|------|---------|----------|------|-------------|-----------|-------|------|-------------|--------|-------------|
| | [1] | [39] | [49] | [14] | [27] | [7] | [3] | [57] | [33] | [48] | [9] | [62] | [2] | [47] | | |
| Prec. (%) | 53.3 | 56.5 | 66.1 | 61.9 | 69.4 | 64.8 | 68.7 | 70.5 | 66.4 | - | 70.4 | 70.0 | 68.8 | 80.0 | 72.7 | 73.1 |
| N. Prec. (%) | 66.3 | 70.5 | 77.8 | 74.6 | 80.0 | 77.1 | 80.1 | 80.0 | 76.8 | 82.2 | 81.6 | 80.9 | 80.0 | 85.4 | 82.9 | 83.3 |
| Success (%) | 57.1 | 60.6 | 71.2 | 66.9 | 73.3 | 70.3 | 74.0 | 75.4 | 72.8 | 75.7 | 75.8 | 75.2 | 74.0 | 81.2 | 78.1 | 78.4 |

Table 4. Comparison results on the GOT-10k test set [23] in terms of average overlap (AO), and success rates (SR) at overlap thresholds 0.5 and 0.75. We show the tracking results without (w/o) and with (w) additional training data (LTC: LaSOT, TrackingNet, and COCO).

| | SiamFC | SiamFCv2 | SiamRPN | SPM | ATOM | DiMP-50 | SiamFC++ | D3S | PrDiMP-50 | DCFST | KYS | Siam-RCNN | TrSiam | | TrDiMP | |
|------------------------|--------|----------|---------|------|------|---------|----------|------|-----------|-------|------|-----------|-------------|--------|-------------|--------|
| | [1] | [45] | [28] | [49] | [7] | [3] | [57] | [33] | [9] | [62] | [2] | [47] | w/o LTC | w/ LTC | w/o LTC | w/ LTC |
| SR _{0.5} (%) | 35.3 | 40.4 | 54.9 | 59.3 | 63.4 | 71.7 | 69.5 | 67.6 | 73.8 | 75.3 | 75.1 | - | 76.6 | 78.7 | 77.7 | 80.5 |
| SR _{0.75} (%) | 9.8 | 14.4 | 25.3 | 35.9 | 40.2 | 49.2 | 47.9 | 46.2 | 54.3 | 49.8 | 51.5 | - | 57.1 | 58.6 | 58.3 | 59.7 |
| AO (%) | 34.8 | 37.4 | 46.3 | 51.3 | 55.6 | 61.1 | 59.5 | 59.7 | 63.4 | 63.8 | 63.6 | 64.9 | 66.0 | 67.3 | 67.1 | 68.8 |

baselines has been largely narrowed (from 4.7% to 1.5% in AO), which reveals the strong tracking potential of a simple pipeline by adequately exploring the temporal information. **Structure Modifications.** Finally, we discuss some architecture details of our transformer: (1) *Shared-weight Self-attention.* Since our transformer is separated into two parallel Siamese tracking branches, the performance obviously drops without the weight-sharing mechanism as shown in Table 2. Due to this weight-sharing design, we also do not stack multiple encoder/decoder layers like the classic transformer [46], which will divide the template and search representations into different feature subspaces. (2) *Feed-forward Network.* Feed-forward network is a basic block in the classic transformer [46], which consists of two heavy-weight fully-connected layers. In the tracking scenario, we observe that this block potentially causes the overfitting issue due to its overmany parameters, which does not bring performance gains and hurts the efficiency. (3) *Head Number.* Classic transformer adopts multi-head attentions (e.g., 8 heads) to learn diverse representations [46]. In the experiments, we observe that increasing the head number slightly improves the accuracy but hinders the tracking efficiency from real-time. We thus choose the single-head attention to achieve a good balance of performance and efficiency.

5.3. State-of-the-art Comparisons

We compare our proposed TrSiam and TrDiMP trackers with the recent state-of-the-art trackers on seven tracking benchmarks including TrackingNet [38], GOT-10k [23], LaSOT [13], VOT2018 [26], Need for Speed [24], UAV123 [36], and OTB-2015 [56].

TrackingNet [38]. TrackingNet is a recently released large-scale benchmark. We evaluate our methods on the test set of TrackingNet, which consists of 511 videos. In this benchmark, we compare our approaches with the state-of-the-art trackers such as DiMP-50 [3], D3S [33], SiamFC++ [57], Retain-MAML [48], DCFST [62], PrDiMP-50 [9], KYS [2], and Siam-RCNN [47]. As shown in Table 3, the proposed TrDiMP achieves a normalized precision score of




Figure 7. Precision and success plots on the LaSOT test set [13].

83.3% and a success score of 78.4%, surpassing previous state-of-the-art trackers such as PrDiMP-50 and KYS. Note that PrDiMP and KYS improve the DiMP tracker via probabilistic regression and tracking scene exploration, representing the current leading algorithms on several datasets. With our designed transformer, the simple Siamese matching baseline (i.e., TrSiam) also shows outstanding performance with a normalized precision score of 82.9% and a success score of 78.1%.

GOT-10k [23]. GOT-10k is a large-scale dataset including more than 10,000 videos. We test our methods on the test set of GOT-10k with 180 sequences. The main characteristic of GOT-10k is that the test set does not have overlap in object classes with the train set, which is designed to assess the generalization of the visual tracker. Following the test protocol of GOT-10k, we further train our trackers with only the GOT-10k training set. As shown in Table 4, in a fair comparison scenario (i.e., without additional training data), both our TrDiMP and TrSiam still outperform other top-performing trackers such as SiamR-CNN [47], DCFST [62], and KYS [2], verifying the strong generalization of our methods to unseen objects.

LaSOT [13]. LaSOT is a recent large-scale tracking benchmark consisting of 1200 videos. The average video length of this benchmark is about 2500 frames, which is more challenging than the previous short-term tracking datasets. Therefore, how to cope with the drastic target appearance changes using temporal context is vital in this dataset. We evaluate our approaches on the LaSOT test set with 280

Table 5. State-of-the-art comparison on the NfS [24], UAV123 [36], and OTB-2015 [56] datasets in terms of AUC score. Both our TrDiMP and TrSiam exhibit outstanding results on all benchmarks with competitive efficiency.

| | KCF [20] | SiamFC [1] | CFNet [45] | MDNet [39] | C-COT [11] | ECO [8] | ATOM [7] | UPDT [4] | SiamRPN++ [27] | DiMP-50 [3] | SiamR-CNN [47] | PrDiMP-50 [9] | DCFST [62] | KYS [2] | TrSiam | TrDiMP |
|---------------|------------|------------|------------|------------|------------|---------|----------|----------|----------------|-------------|----------------|---------------|-------------|---------|-------------|-------------|
| NfS [24] | 21.7 | - | - | 42.9 | 48.8 | 46.6 | 58.4 | 53.7 | 50.2 | 62.0 | 63.9 | 63.5 | 64.1 | 63.5 | 65.8 | 66.5 |
| UAV123 [36] | 33.1 | 49.8 | 43.6 | 52.8 | 51.3 | 52.2 | 64.2 | 54.5 | 61.3 | 65.3 | 64.9 | 68.0 | - | - | 67.4 | 67.5 |
| OTB-2015 [56] | 47.5 | 58.2 | 56.8 | 67.8 | 68.2 | 69.1 | 66.9 | 70.2 | 69.6 | 68.4 | 70.1 | 69.6 | 70.9 | 69.5 | 70.8 | 71.1 |
| Speed (FPS) | 270 | 86 | 75 | 1 | 0.3 | 8 | 35 | <1 | 30 | 35 | 4.7 | 30 | 25 | 20 | 35 | 26 |




Figure 8. Expected average overlap (EAO) graph with trackers ranked from right to left. Our TrDiMP and TrSiam outperform all the participant trackers on the VOT2018 [26].

videos. The precision and success plots of the state-of-the-art methods are shown in Figure 7, where the recently proposed C-RPN [14], SiamRPN++ [27], ATOM [7], DiMP-50 [3], and PrDiMP-50 [9] are included for comparison. Our TrSiam and TrDiMP outperform aforementioned methods by a considerable margin. To the best of our knowledge, SiamR-CNN [47] achieves the current best result on the LaSOT. Overall, our TrDiMP (63.9% AUC and 26 FPS) exhibits very competitive performance and efficiency in comparison with SiamR-CNN (64.8% AUC and 4.7 FPS).

VOT2018 [26]. VOT2018 benchmark contains 60 challenging videos. The performance on this dataset is evaluated using the expected average overlap (EAO), which takes both accuracy (average overlap over successful frames) and robustness (failure rate) into account. As shown in Figure 8, our TrSiam and TrDiMP clearly outperform all the participant trackers on the VOT2018.

In Table 6, we further show the accuracy, robustness, and EAO scores of the recent top-performing trackers including SiamRPN++ [27], DiMP-50 [3], PrDiMP-50 [9], Retain-MAML [48], KYS [2], and D3S [33]. Compared with these recently proposed approaches, our TrDiMP approach still exhibits satisfactory results. Among all the compared trackers, only D3S slightly outperforms our TrDiMP, which is trained using additional data with segmentation annotations for accurate mask prediction.

NfS [24]. NfS dataset contains 100 challenging videos with fast-moving objects. We evaluate our TrSiam and TrDiMP on the 30 FPS version of NfS. The AUC scores of comparison approaches are shown in Table 5. Our approaches set new state-of-the-art records on this benchmark. The proposed TrDiMP surpasses previous top-performing trackers such as DCFST [62] and SiamR-CNN [47]. Note that the recent SiamR-CNN utilizes a powerful ResNet-101 for ob-

Table 6. Comparison with recent state-of-the-art trackers on the VOT2018 [26] in terms of accuracy (A), robustness (R), and expected average overlap (EAO).

| | SiamRPN++ [27] | DiMP-50 [3] | PrDiMP-50 [9] | Retain-MAML [48] | KYS [2] | D3S [33] | TrDiMP |
|---------|----------------|-------------|---------------|------------------|--------------|--------------|--------------|
| A (↑) | 0.600 | 0.597 | 0.618 | 0.604 | 0.609 | 0.640 | 0.600 |
| R (↓) | 0.234 | 0.153 | 0.165 | 0.159 | 0.143 | 0.150 | 0.141 |
| EAO (↑) | 0.414 | 0.440 | 0.442 | 0.452 | 0.462 | 0.489 | 0.462 |

ject re-detection. Our simple TrSiam, without sophisticated models or online optimization techniques, still outperforms existing methods and operates in real-time.

UAV123 [36]. This benchmark includes 123 aerial videos collected by the low-attitude UAV platform. The proposed trackers also achieve promising results in comparison to the recent remarkable approaches in Table 5. Specifically, our TrDiMP performs on par with PrDiMP-50 [9], which represents the current best algorithm on this benchmark.

OTB-2015 [56]. OTB-2015 is a popular tracking benchmark with 100 challenging videos. As shown in Table 5, on this dataset, our TrDiMP achieves an AUC score of 71.1%, surpassing the recently proposed SiamRPN++ [27], PrDiMP-50 [9], SiamR-CNN [47], and KYS [2]. With the proposed transformer, our Siamese matching based TrSiam also performs favorably against existing state-of-the-art approaches with an AUC score of 70.8%.

6. Conclusion

In this work, we introduce the transformer structure to the tracking frameworks, which bridges the isolated frames in the video flow and conveys the rich temporal cues across frames. We show that by carefully modifying the classic transformer architecture, it favorably suits the tracking scenario. With the proposed transformer, two popular trackers gain consistent performance improvements and set several new state-of-the-art records on prevalent tracking datasets. To our best knowledge, this is the first attempt to exploit the transformer in the tracking community, which preliminarily unveils the tracking potential hidden in the frame-wise relationship. In the future, we intend to further explore the rich temporal information among individual video frames.

Acknowledgements. This work was supported in part by the National Natural Science Foundation of China under Contract 61836011, 61822208, and 61836006, and in part by the Youth Innovation Promotion Association CAS under Grant 2018497. It was also supported by the GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC.

References

- [1] Luca Bertinetto, Jack Valmadre, João F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *ECCV Workshops*, 2016.
- [2] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Know your surroundings: Exploiting scene information for object tracking. In *ECCV*, 2020.
- [3] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *ICCV*, 2019.
- [4] Goutam Bhat, Joakim Johnander, Martin Danelljan, Fahad Shahbaz Khan, and Michael Felsberg. Unveiling the power of deep tracking. In *ECCV*, 2018.
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E Hinton. A simple framework for contrastive learning of visual representations. *arXiv: 2002.05709*, 2020.
- [7] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Atom: Accurate tracking by overlap maximization. In *CVPR*, 2019.
- [8] Martin Danelljan, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Eco: Efficient convolution operators for tracking. In *CVPR*, 2017.
- [9] Martin Danelljan, Luc Van Gool, and Radu Timofte. Probabilistic regression for visual tracking. In *CVPR*, 2020.
- [10] Martin Danelljan, Gustav Häger, Fahad Khan, and Michael Felsberg. Accurate scale estimation for robust visual tracking. In *BMVC*, 2014.
- [11] Martin Danelljan, Andreas Robinson, Fahad Shahbaz Khan, and Michael Felsberg. Beyond correlation filters: Learning continuous convolution operators for visual tracking. In *ECCV*, 2016.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [13] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. Lasot: A high-quality benchmark for large-scale single object tracking. In *CVPR*, 2019.
- [14] Heng Fan and Haibin Ling. Siamese cascaded region proposal networks for real-time visual tracking. In *CVPR*, 2019.
- [15] Hamed Kiani Galoogahi, Ashton Fagg, and Simon Lucey. Learning background-aware correlation filters for visual tracking. In *ICCV*, 2017.
- [16] Junyu Gao, Tianzhu Zhang, and Changsheng Xu. Graph convolutional tracking. In *CVPR*, 2019.
- [17] Qing Guo, Wei Feng, Ce Zhou, Rui Huang, Liang Wan, and Song Wang. Learning dynamic siamese network for visual object tracking. In *ICCV*, 2017.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [19] David Held, Sebastian Thrun, and Silvio Savarese. Learning to track at 100 fps with deep regression networks. In *ECCV*, 2016.
- [20] João F Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista. High-speed tracking with kernelized correlation filters. *TPAMI*, 37(3):583–596, 2015.
- [21] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [22] Chen Huang, Simon Lucey, and Deva Ramanan. Learning policies for adaptive tracking with deep feature cascades. In *ICCV*, 2017.
- [23] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *arXiv preprint arXiv:1810.11981*, 2018.
- [24] Hamed Kiani Galoogahi, Ashton Fagg, Chen Huang, Deva Ramanan, and Simon Lucey. Need for speed: A benchmark for higher frame rate object tracking. In *ICCV*, 2017.
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [26] Matej Kristan, Ales Leonardis, Jiri Matas, Michael Felsberg, Roman Pflugfelder, Luka Cehovin Zajc, Tomas Vojir, Goutam Bhat, Alan Lukezic, Abdelrahman Eldesokey, et al. The sixth visual object tracking vot2018 challenge results. In *ECCV Workshops*, 2018.
- [27] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *CVPR*, 2019.
- [28] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *CVPR*, 2018.
- [29] Feng Li, Cheng Tian, Wangmeng Zuo, Lei Zhang, and Ming-Hsuan Yang. Learning spatial-temporal regularized correlation filters for visual tracking. In *CVPR*, 2018.
- [30] Xin Li, Chao Ma, Baoyuan Wu, Zhenyu He, and Ming-Hsuan Yang. Target-aware deep tracking. In *CVPR*, 2019.
- [31] 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 *ECCV*, 2014.
- [32] Xiankai Lu, Chao Ma, Bingbing Ni, Xiaokang Yang, Ian Reid, and Ming-Hsuan Yang. Deep regression tracking with shrinkage loss. In *ECCV*, 2018.
- [33] Alan Lukezic, Jiri Matas, and Matej Kristan. D3s-a discriminative single shot segmentation tracker. In *CVPR*, 2020.
- [34] Alan Lukezic, Tomas Vojir, Luka Cehovin Zajc, Jiri Matas, and Matej Kristan. Discriminative correlation filter with channel and spatial reliability. In *CVPR*, 2017.
- [35] Chao Ma, Jia-Bin Huang, Xiaokang Yang, and Ming-Hsuan Yang. Hierarchical convolutional features for visual tracking. In *ICCV*, 2015.
- [36] Matthias Mueller, Neil Smith, and Bernard Ghanem. A benchmark and simulator for uav tracking. In *ECCV*, 2016.
- [37] Matthias Mueller, Neil Smith, and Bernard Ghanem. Context-aware correlation filter tracking. In *CVPR*, 2017.

- [38] Matthias Müller, Adel Bibi, Silvio Giancola, Salman Al-Subaihi, and Bernard Ghanem. Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In *ECCV*, 2018.
- [39] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016.
- [40] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *ICML*, 2018.
- [41] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.
- [42] Yibing Song, Chao Ma, Lijun Gong, Jiawei Zhang, Rynson Lau, and Ming-Hsuan Yang. Crest: Convolutional residual learning for visual tracking. In *ICCV*, 2017.
- [43] Gabriel Synnaeve, Qiantong Xu, Jacob Kahn, Edouard Grave, Tatiana Likhomanenko, Vineel Pratap, Anuroop Sriram, Vitaliy Liptchinsky, and Ronan Collobert. End-to-end asr: from supervised to semi-supervised learning with modern architectures. *arXiv preprint arXiv:1911.08460*, 2019.
- [44] Ran Tao, Efstratios Gavves, and Arnold WM Smeulders. Siamese instance search for tracking. In *CVPR*, 2016.
- [45] Jack Valmadre, Luca Bertinetto, João F Henriques, Andrea Vedaldi, and Philip HS Torr. End-to-end representation learning for correlation filter based tracking. In *CVPR*, 2017.
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [47] Paul Voigtlaender, Jonathon Luiten, Philip H. S. Torr, and Bastian Leibe. Siam r-cnn: Visual tracking by re-detection. In *CVPR*, 2020.
- [48] Guangting Wang, Chong Luo, Xiaoyan Sun, Zhiwei Xiong, and Wenjun Zeng. Tracking by instance detection: A meta-learning approach. In *CVPR*, 2020.
- [49] Guangting Wang, Chong Luo, Zhiwei Xiong, and Wenjun Zeng. Spm-tracker: Series-parallel matching for real-time visual object tracking. In *CVPR*, 2019.
- [50] Ning Wang, Yibing Song, Chao Ma, Wengang Zhou, Wei Liu, and Houqiang Li. Unsupervised deep tracking. In *CVPR*, 2019.
- [51] Ning Wang, Wengang Zhou, Guojun Qi, and Houqiang Li. Post: Policy-based switch tracking. In *AAAI*, 2020.
- [52] Ning Wang, Wengang Zhou, Yibing Song, Chao Ma, Wei Liu, and Houqiang Li. Unsupervised deep representation learning for real-time tracking. *IJCV*, 129(2):400–418, 2021.
- [53] Ning Wang, Wengang Zhou, Qi Tian, Richang Hong, Meng Wang, and Houqiang Li. Multi-cue correlation filters for robust visual tracking. In *CVPR*, 2018.
- [54] Qiang Wang, Zhu Teng, Junliang Xing, Jin Gao, Weiming Hu, and Stephen Maybank. Learning attentions: Residual attentional siamese network for high performance online visual tracking. In *CVPR*, 2018.
- [55] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [56] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Object tracking benchmark. *TPAMI*, 37(9):1834–1848, 2015.
- [57] Yinda Xu, Zeyu Wang, Zuoxin Li, Yuan Ye, and Gang Yu. Siamfc++: Towards robust and accurate visual tracking with target estimation guidelines. In *AAAI*, 2020.
- [58] Tianyu Yang and Antoni B Chan. Learning dynamic memory networks for object tracking. In *ECCV*, 2018.
- [59] Yuechen Yu, Yilei Xiong, Weilin Huang, and Matthew R. Scott. Deformable siamese attention networks for visual object tracking. In *CVPR*, 2020.
- [60] Lichao Zhang, Abel Gonzalez-Garcia, Joost van de Weijer, Martin Danelljan, and Fahad Shahbaz Khan. Learning the model update for siamese trackers. In *ICCV*, 2019.
- [61] Zhipeng Zhang and Houwen Peng. Deeper and wider siamese networks for real-time visual tracking. In *CVPR*, 2019.
- [62] Linyu Zheng, Ming Tang, Yingying Chen, Jinqiao Wang, and Hanqing Lu. Learning feature embeddings for discriminant model based tracking. In *ECCV*, 2020.
- [63] Zheng Zhu, Qiang Wang, Bo Li, Wei Wu, Junjie Yan, and Weiming Hu. Distractor-aware siamese networks for visual object tracking. In *ECCV*, 2018.
- [64] Zheng Zhu, Wei Wu, Wei Zou, and Junjie Yan. End-to-end flow correlation tracking with spatial-temporal attention. In *CVPR*, 2018.