






Diff-Tracker: Text-to-Image Diffusion Models are Unsupervised Trackers

Zhengbo Zhang¹, Li Xu¹, Duo Peng¹,
Hossein Rahmani², and Jun Liu^{1,2} *

¹ Singapore University of Technology and Design
{zhengbo_zhang, li_xu, duo_peng}@mymail.sutd.edu.sg

² Lancaster University
h.rahmani@lancaster.ac.uk, j.liu81@lancaster.ac.uk

Abstract. We introduce Diff-Tracker, a novel approach for the challenging unsupervised visual tracking task leveraging the pre-trained text-to-image diffusion model. Our main idea is to leverage the rich knowledge encapsulated within the pre-trained diffusion model, such as the understanding of image semantics and structural information, to address unsupervised visual tracking. To this end, we design an initial prompt learner to enable the diffusion model to recognize the tracking target by learning a prompt representing the target. Furthermore, to facilitate dynamic adaptation of the prompt to the target’s movements, we propose an online prompt updater. Extensive experiments on five benchmark datasets demonstrate the effectiveness of our proposed method, which also achieves state-of-the-art performance.

Keywords: Visual object tracking · Text-to-image diffusion model · Unsupervised learning

1 Introduction

Visual object tracking constitutes a core task in the field of computer vision, finding extensive applications in areas ranging from autonomous driving [4, 16] to robotics [3, 38]. In the recent development in visual object tracking, deep learning based trackers [7, 10, 33, 36, 59, 60, 62, 63] have emerged as a prevailing paradigm. These trackers exhibit a strong reliance on data, necessitating a substantial volume of annotated data for supervised training. Owing to the high cost and time demands associated with manual data annotation, unsupervised visual tracking has recently experienced a surge in attention. Although previous researchers have made significant efforts [46, 51, 68], unsupervised object tracking remains a substantial challenge due to the difficulty in effectively exploiting the rich semantic and structural information of video frames [47, 51], as well as the abundant contextual relationships within videos [68].

*corresponding author

At the same time, the pre-trained text-to-image diffusion model has achieved exemplary performance across various realms, including text-to-image generation [21, 42] and text-to-video generation [26]. For instance, in the text-to-image area, the pre-trained text-to-image diffusion model, such as Stable Diffusion [42], has demonstrated remarkable capabilities in generating images that are not only diverse but also rich in detail and adhere to reasonable spatial structures, all in response to user-defined prompts. Such impressive results imply that the pre-trained text-to-image diffusion model exhibits an extensive understanding of visual representations, ranging from pixel-level semantic details to spatial layouts, including aspects like object texture, object shape, and spatial arrangement within the image. Additionally, researchers [49] verify that the pre-trained diffusion model, despite not being trained or fine-tuned on video data, exhibits commendable capabilities in understanding video contextual relationships. Given that the pre-trained text-to-image diffusion model demonstrates the understanding (knowledge) of the semantic and structural information of video frames, and contextual relationships within video, a natural question arises: *Can we leverage the rich knowledge encapsulated in the pre-trained text-to-image diffusion model for performing unsupervised visual tracking?*

Nevertheless, it is non-trivial to leverage the knowledge implicitly embedded in the pre-trained text-to-image diffusion model for unsupervised object tracking in videos. Such models are engineered for image generation from text prompts and are not inherently suited for visual object tracking. To handle this problem, we delve into the underlying characteristic of the text-to-image diffusion model. **Firstly**, we revisit the input type (text prompts) and output type (images) of the text-to-image diffusion model, interpreting the model’s function from a novel perspective. We view the text-to-image diffusion model as a bridge that connects the semantic information of the input prompts with the content of the output images. **Secondly**, this semantic connection between text prompts and output images is specifically manifested in the cross-attention maps within the text-to-image diffusion model [27]. The text prompts are capable of activating areas on the cross-attention maps that are semantically related to the prompts. In this paper, activated regions mean that the pixel values in these areas on the attention maps are high. This suggests that if we learn a prompt for the tracking target, the prompt can enable the text-to-image diffusion model to activate region of the target object on the cross-attention maps of input video frames. Consequently, we can leverage the rich knowledge embedded within the pre-trained text-to-image diffusion model to perform the challenging unsupervised object tracking.

However, learning a prompt for the target object is not straightforward: i) Since the relationship between the target and the background is beneficial for tracking the target in scenarios where it undergoes significant appearance deformations or occlusions [14, 69], it presents a challenge that the learned prompt should encode rich relationships between the target object and backgrounds. ii) The continuous motion of the target object leads to changes in the appearance of the object, as well as the relationships between the target object and the

background. This poses another challenge of how to update the prompt online to adapt to these changes.

To address these challenges, we propose Diff-Tracker, which consists of an initial prompt learner and an online prompt updater. The initial prompt learner is devised to learn an initial prompt for tracking the target object. To be exact, through this module, we learn the initial prompt that can accurately activate the target object’s region in the cross-attention map of the first frame. To encapsulate more information about the relationship between the target object and the background, we design an attention harmonization method to combine the original cross-attention map and cross-attention map enhanced by the self-attention map from the text-to-image diffusion model.

The designed online prompt updater enables the learned prompt to update according to the target’s motion. To capture the real-time motion of the target object, a motion information extractor is proposed to extract target-conditioned motion information between the current frame and the previous one. However, relying solely on the extraction of motion information between two consecutive frames may lead to unreliable results. This is because phenomena like occlusions and illumination changes can affect the accuracy of short-term motion information, potentially disrupting spatio-temporal continuity.

To address this issue, in addition to the short-term motion information, we also incorporate long-term motion information, thereby enhancing robustness in maintaining spatio-temporal continuity [6].

We summarize our contributions: i) From a novel perspective, we perform unsupervised object tracking, by leveraging the rich knowledge embedded in the pre-trained text-to-image diffusion model. ii) To harness the potential of the diffusion models for unsupervised object tracking, we introduce Diff-Tracker, a novel framework with crafted design components. The components include an initial prompt learner for learning an initial prompt representing the target and an online prompt updater for updating the learned prompt based on the target’s motion information. iii) Our method achieves state-of-the-art performance in unsupervised object tracking task.

2 Related Work

Unsupervised visual tracking. Thanks to deep learning techniques, significant advancements have been made in various domains of computer vision, such as object detection [22, 48, 55], semantic segmentation [57, 61, 64], image classification [65–67], video understanding [5, 19, 52, 56], and unsupervised visual tracking [1, 46, 47, 50, 68]. The groundbreaking deep learning-based unsupervised visual tracking work UDT [50] develops a tracker based on Discriminative Correlation Filters. This tracker is trained through forward-backward tracking of frames, under the guidance of a consistency loss function. Besides, s^2 siamfc [47] utilizes a Siamese pipeline, similar to SiamFC [1], for training a foreground/background classifier with single-frame pairs. This method incorporates learning adversarial masking techniques to generate template-search pairs from identity frame. Fol-

lowing s^2 siamfc, state-of-the-art unsupervised trackers [46, 68] also employ similar siamese network structure. Differently, here we leverage the rich knowledge embedded within text-to-image diffusion models for the first time to perform unsupervised visual tracking.

Text-to-image diffusion models. Recent developments [21, 42, 45] in text-to-image diffusion models have significantly enhanced the capacity of generating diverse images from input prompts. With the advent of these pre-trained text-to-image diffusion models, *e.g.* Stable Diffusion [42] and Imagen [45], there has been a surge in research aiming to leverage powerful capabilities of these models to further advance developments in other fields, including image editing [25], text-to-video generation [26], domain adaptation [40], object counting [24], pose estimation [17, 41], image inpainting [54], human mesh recovery [15], single-view object reconstruction [35], and semantic segmentation [27]. The landscape of these fields is undergoing significant transformation, driven by the versatility and effectiveness of the diffusion technique. Different from their work, we extend the application of the text-to-image diffusion models to the domain of unsupervised visual tracking.

3 Preliminaries: Text-to-Image Diffusion Models

Text-to-image diffusion models are designed to reconstruct an image from a random Gaussian noise. This is achieved by progressively denoising the noise through a reverse diffusion process, which is conditioned on a text prompt p . Below, we use the Stable Diffusion [42], a powerful diffusion model, as an example to illustrate the training process of the text-to-image diffusion models, as well as the cross-attention and self-attention mechanisms within the diffusion model.

Training process. During the training phase, an input image I and a corresponding text prompt p are encoded by an image encoder \mathcal{E}_i and a text encoder \mathcal{E}_t , respectively. Noise ϵ is added to the latent representation of the input image $\mathcal{E}_i(I)$, which is used for denoising training. The training objective of the Stable Diffusion model is to predict the added noise ϵ in the encoded image. This is mathematically formulated as:

$$L_{DM} = \mathbb{E}_{\mathcal{E}_i(I), \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_\theta(z_t, t, \mathcal{E}_t(p))\|_2^2], \quad (1)$$

Where t denotes the time step in the diffusion process, and ϵ_θ is a denoising UNet [43] network.

Cross-attention layers. In the Stable Diffusion, interplay between the image content and text prompt is encapsulated within the cross-attention layers of the denoising UNet. Specifically, within a cross-attention layer, the query Q_c

originates from the noisy image latent, while the key K_c and value V_c are derived from the text embedding $\mathcal{E}_t(p)$. The cross-attention map can be formulated as:

$$M_c = \text{Softmax} \left(\frac{Q_c K_c^T}{\sqrt{d}} \right), \quad (2)$$

where d represents the projection dimension of Q_c , K_c , and V_c . The cross-attention map, M_c , obtained by calculating the correlation between the image content represented in Q_c and the semantic text content in K_c , serves to effectively highlight the image regions semantically related to the text prompt.

Self-attention layers. In the architecture of the Stable Diffusion, along with the cross-attention layers, self-attention layers are also integrated within the UNet. These layers are instrumental in capturing semantic correlations among pixels in an image. Specifically, the self-attention map M_s can be defined as:

$$M_s = \text{Softmax} \left(\frac{Q_s K_s^T}{\sqrt{d}} \right), \quad (3)$$

which is in a same format with the cross-attention map M_c . The only difference is that, the query Q_s , the key K_s , and the value V_s in the self-attention map M_s are all derived from the noisy image latent representation. The self-attention map M_s encapsulates information about the semantic relationships among pixels. This is because the derivation of the self-attention map involves calculating the correlation between Q_s and K_s , both of which originate from the image pixels.

4 Diff-Tracker

In this section, we first revisit the task setup of the unsupervised visual tracking task and briefly introduce framework of our Diff-Tracker (Sec. 4.1), followed by a detailed description of two components (*i.e.* the initial prompt learner (Sec. 4.2) and online prompt updater (Sec. 4.3)) of the Diff-Tracker.

4.1 Task Definition and Our Framework

Task definition. Given the location (bounding box annotation) of an arbitrary target in the first frame, unsupervised visual tracking task aims to accurately predict the locations of this target within the following video frame sequence, while the tracker is trained exclusively on unlabeled videos.

Our framework. The purpose of designing the Diff-Tracker is to utilize the abundant knowledge embedded in the pre-trained text-to-image diffusion model to assist in unsupervised visual tracking. Inspired by [27], we observe that inputting an image and a text prompt into a pre-trained text-to-image diffusion model allows the model to activate regions semantically related to the text

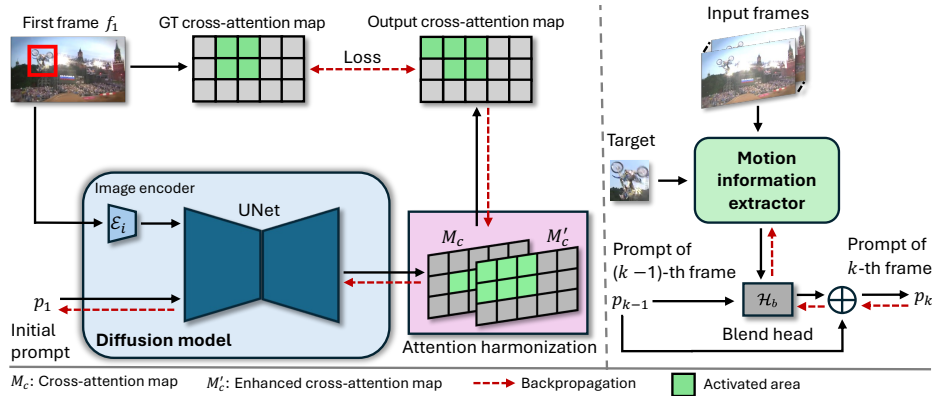


Fig. 1: The framework of the Diff-Tracker consists of the initial prompt learner on the left side of the figure and the online prompt updater on the right. The prompt updated through the online prompt learner is input into the network of initial prompt updater to obtain the output cross-attention map. This map is used to compute the loss for updating the online prompt updater by comparison with the GT cross-attention map.

prompt on the cross-attention map of the UNet, leveraging the extensive knowledge embedded within the diffusion model. Furthermore, this prompt can accurately activate regions representing its semantics across different input images, demonstrating strong generalization capabilities.

Based on this, we aim to learn a prompt representing the target to facilitate unsupervised visual tracking with the aid of the diffusion model. Since we do not have text and some targets cannot be explained with text, in this paper, we use prompt embedding as the prompt that needs to be learned. However, acquiring such a prompt is not straightforward. One challenge lies in encoding the relationships between the target and the background, which are crucial for effective target tracking [69]. Another challenge pertains to updating the learned prompt based on the target’s motion, considering that the target’s appearance may undergo changes as it moves.

To address these issues, we divide the process of learning a prompt representing the target into two distinct phases. As shown in Fig. 1, the first phase involves learning an initial prompt activating the region corresponding to the target object in the first frame, for which we design an initial prompt learner (Sec. 4.2). In the second phase, the learned initial prompt undergoes updates to accommodate the target’s motion, ensuring continuous tracking of the target throughout the video sequence. To achieve this objective, we introduce an online prompt updater (Sec. 4.3). By integrating the functionalities of the initial prompt learner and the online prompt updater, the Diff-Tracker framework is established, enabling effective and continuous tracking of the target.

4.2 Initial Prompt Learner

The goal of the initial prompt learner is to make the text-to-image diffusion model aware of the target to be tracked. Inspired by [27], we achieve this goal by learning a prompt that activates the region corresponding to the target on the cross-attention map within the diffusion model. Besides, in the process of visual tracking, we may encounter common perturbations such as significant deformation or occlusion of the target, as well as the presence of nearby objects with appearances similar to the target. Under these circumstances, the relationship between the target and its background serves as crucial information for tracking [14, 69].

To encapsulate the relationship between the target and the background in both the initial prompt and the updated prompt, we design an attention harmonization method that utilizes the self-attention mechanism of the diffusion model, because the self-attention map encodes semantic relationships between pixels. It is important to highlight that our attention harmonization method further allows the prompt, updated by the online prompt updater, to encapsulate the relationship between the target and its background in subsequent frames. This capability stems from the online prompt updater’s utilization of the model from the initial prompt learner and its attention harmonization method.

The subsequent subsections delineate the designed initial prompt learner.

Attention harmonization. Considering the self-attention map of the pre-trained text-to-image diffusion model is able to capture semantic relationships between pixels (as mentioned in the Section 3), we develop a harmonization method (see Fig. 1) to combine the self-attention map with the cross-attention map, aiming to encode the relationship between the target and its background within the learned prompt. Specifically, for the cross-attention map M_c extracted from the UNet (Eq. 2), each pixel on M_c can be associated with a self-attention map M_s (Eq. 3) to reflect the pixel’s relationship with other pixels. We enhance the cross-attention map M_c using the self-attention map M_s , thereby encoding the relationships among pixels. The enhanced cross-attention map M'_c is calculated as follows:

$$M'_c(:, :) = \sum_{i=1} \sum_{j=1} M_c(i, j) \cdot M_s(i, j, :, :), \quad (4)$$

where $M_c(i, j)$ represents the attention value at position (i, j) on the cross-attention map M_c , while $M_s(i, j, :, :)$ denotes the self-attention map related to pixel (i, j) of the cross-attention map M_c .

Upon obtaining M'_c , which encodes the inter-pixel relationships, we proceed to integrate M'_c with the original cross-attention map M_c to form the final output cross-attention map \mathcal{M} . To achieve this aim, we resize M'_c to the same size as M_c and perform an element-wise weighted summation of these two attention maps to obtain the final output cross-attention map \mathcal{M} . The map \mathcal{M} can be defined as follows:

$$\mathcal{M} = (1 - \alpha) \cdot M'_c + \alpha \cdot M_c, \quad (5)$$

where α represents a hyper parameter to balance the weights of M'_c and M_c .

We optimize the prompt embedding of the initial prompt by calculating the MSE (Mean Squared Error) loss between \mathcal{M} and the GT cross-attention map. In this paper, the GT cross-attention map is defined as an attention map that only activates the bounding box area.

Learning of initial prompt. Finally, we describe the learning process of initial prompt p_1 for target object in the first frame f_1 . Specifically, for the given f_1 , we firstly employ the diffusion model’s image encoder \mathcal{E}_i to encode f_1 into latent space, and get encoded representation of f_1 . Following the workflow of the diffusion model, noise is added to this encoded representation, resulting in a noisy latent representation z of f_1 . z is fed into the denoising UNet network ϵ_θ , accompanied by the prompt p_1 . Subsequently, we extract the cross-attention map M_c from the UNet network ϵ_θ and enhance it using the self-attention map M_s . Both the original cross-attention map M_c and the enhanced one M'_c are then fused through our designed attention harmonization method to obtain the integrated cross-attention map \mathcal{M} . Ultimately, we calculate the normalized MSE loss between \mathcal{M} and the GT cross-attention map \mathcal{F}_1 . Overall, the loss function is defined as follows:

$$L = \|\mathcal{M} - \mathcal{F}_1\|_2^2 + L_{DM}, \quad (6)$$

where L_{DM} is the loss of the diffusion model (see the Eq. 1). The use of loss L_{DM} is to restrict the learned prompt p_1 resides within the text embedding space understandable by the diffusion model. It is worth noting that during the process of learning p_1 , the parameters of the diffusion model are frozen.

The learned initial prompt p_1 is subsequently fed into the online prompt updater for online updating of the prompt.

4.3 Online Prompt Updater

Now that we have obtained the initial prompt p_1 representing the target object from the first frame, we can proceed with visual tracking in the subsequent frames. In an ideal scenario, we can utilize the initial prompt to track the target throughout the subsequent frames. Yet, the object may undergo appearance changes due to prolonged motion over time, and the initial prompt p_1 , might fail to continuously track the object. To address this issue, we design an online prompt updater that takes into account the motion information of the target to dynamically update the prompt for the subsequent frames.

To capture the real-time motion of the target, we utilize the motion information between the current frame and the previous frame as a basis for updating the prompt. However, due to occlusions and changes in illumination during the target’s motion, such target-conditioned short-term motion between two consecutive frames may exhibit limited spatio-temporal coherence. Thus, to perform a more reliable prompt update process, instead of solely relying on short-term

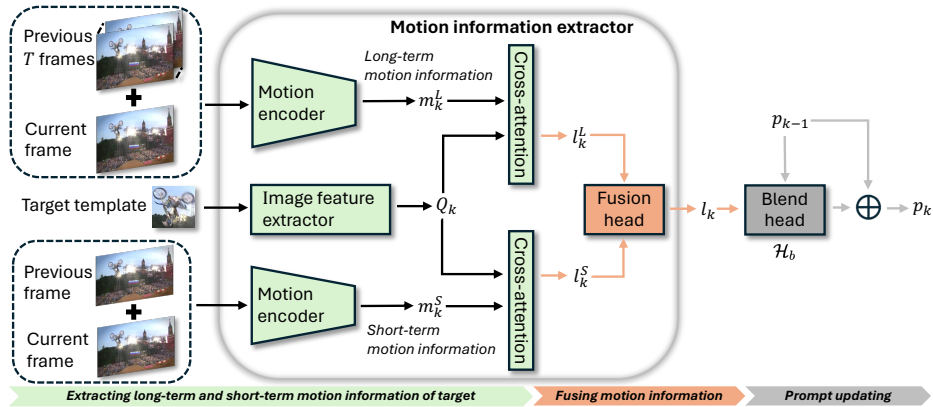


Fig. 2: The detailed architectures of the online prompt updater.

motion to update the prompt, we also incorporate the long-term motion information during the update process. This is because that the long-term motion of moving objects typically demonstrates strong spatio-temporal continuity [6].

Specifically, for each following frame, we first encode the target’s motion information using a motion information extractor. To incorporate the encoded target motion information into the updated prompt p_k , we design a blend head \mathcal{H}_b to combine the encoded motion information with the prompt p_{k-1} of the previous frame. Finally, the output from the designed blend head \mathcal{H}_b is combined with the previous frame’s target prompt p_{k-1} via a residual manner. This residual manner is designed to enhance the stability of our online prompt updater.

We sequentially introduce these steps in the following subsections.

Motion information extractor. To enable reliable updates of the learned prompt based on the target’s motion, we first encode the long-term and short-term motion information. Specifically, as can be seen in Fig. 2, we stack the current frame and previous T frames preceding it as a video sequence and input them into one motion encoder to obtain the long-term motion information m_k^L , and we input the current frame and the previous frame into another motion encoder to obtain the short-term motion information m_k^S .

Since the two types of encoded motion information are based on the entire image and we only track the target, we need to derive the target-conditioned motion information (denoted as l_k below) based on these two types of global motion. To this end, we utilize the cross-attention mechanism to establish connections between the target’s appearance features and the two kinds of motion information. Specifically, we construct one cross-attention map by using the target’s appearance features as the query, with the long-term motion information serving as the key and value. Furthermore, we construct another cross-attention map using the appearance features of the target as the query and the short-term motion information as the key and value. The two cross-attention maps

are illustrated in the following equation:

$$l_k^L = \text{Cross-Attn}(Q_k, m_k^L), \quad l_k^S = \text{Cross-Attn}(Q_k, m_k^S), \quad (7)$$

where l_k^L and l_k^S respectively represent the long-term and short-term motion information conditioned on the target, while Q_k denotes the query. The query Q_k is obtained from the appearance features of the target, which are extracted using an image feature extractor. We input the target template from the first frame into the image feature extractor to obtain the target’s appearance features.

Then, as shown in Fig. 2, to obtain l_k , we design a fusion head to merge the long-term and short-term motion information, l_k^L and l_k^S . The fusion head is a multi-layer perceptron (MLP) network consisting of two fully connected layers and a ReLU activation layer. The obtained target’s motion information l_k serves as crucial information for updating the prompt.

Process of prompt updating. As can be seen in Fig. 2, we update the prompt representing the target in accordance with the target’s motion, accommodating changes in appearance that may occur as a result of such motion. To be specific, we use a blend head \mathcal{H}_b to fuse the target’s motion information l_k , which is obtained through the motion information extractor, with the prompt of the previous frame. Finally, prompt for the current k -th frame is derived by combining the output from the blend head \mathcal{H}_b with the prompt of the $(k-1)$ -th frame p_{k-1} via a residual manner. Formally, the above prompt updating process for the k -th frame is defined as:

$$p_k = (1 - \beta) \cdot \mathcal{H}_b(p_{k-1} + l_k) + \beta \cdot p_{k-1}, \quad (8)$$

where β is a hyper parameter to balance the weight of different terms. The loss for updating the online prompt updater is obtained by calculating the MSE loss between the cross-attention map generated by the prompt p_k and the GT cross-attention map of the k -th frame.

The obtained p_k obtained through the online updating is used for tracking the target on the k -th frame.

5 Experiments

In this section, we first introduce the implementation details (Sec. 5.1) of the Diff-Tracker, as well as the main experimental results (Sec. 5.2) and ablation studies (Sec. 5.3).

5.1 Implementation Details

Pseudo label generation. To obtain pseudo labels, we follow prior work [68] by employing an off-the-shelf optical flow model [34] on datasets GOT-10k [23], ImageNet VID [44], LaSOT [13], and YouTube-VOS [58]. The sampling strategy

for the pseudo labels is in accordance with [68]. In this paper, for fairness in comparison with previous unsupervised trackers [46,68], the GT cross-attention is defined as an attention map that exclusively activates within the area of the bounding box.

Model architecture. Our initial prompt learner is a pre-trained text-to-image diffusion model [42]. In the online prompt updater, both two motion encoders are designed by extending a ResNet18 [18] model, pre-trained on the ImageNet [12] dataset, with two Conv3D layers. The image feature extractor is a also ResNet-18 model pre-trained on the ImageNet dataset. Both the fusion head and blend head are the multi-layer perceptron network consisting of two fully connected layers and a ReLU activation layer.

Training. Our experiments are carried out using RTX 3090 GPUs. For these experiments, we set our input image size to 512×512 , the shape of feature maps (M_c , M'_c , and \mathcal{M}) to $1 \times 64 \times 64$. The dimension of the prompt embedding is 1024. The hyper parameters α and β are set to 0.5 and 0.7, respectively. During training, the parts that need to be learned in our Diff-Tracker are the initial prompt and the parameters in online prompt updater, and the pre-trained text-to-image diffusion model is frozen. For the training of the initial prompt, we set the learning rate at 5×10^{-3} for 3 epochs with the Adam optimizer [28]. Besides, we train the online prompt updater for 35 epochs with the Adam optimizer and learning rate of 5×10^{-4} .

Testing. When testing with a given video, we first utilize the initial prompt learner with the first frame to learn the target’s initial prompt. The experimental setup for learning this initial prompt is the same as the setup for learning the initial prompt during the training phase. Starting from the sixth frame, we update the learned prompt with the online prompt updater for every subsequent frame and use the updated prompt for the visual tracking. To conform to the requirements of standard tracking benchmarks, we follow [39] to report the smallest axis-aligned bounding box that encapsulates the activated area in the cross-attention map as the bounding box.

Evaluation datasets. We follow previous work [46] by conducting experiments on five challenging tracking datasets, including OTB2015 [53], VOT2016 [29], VOT2018 [30], TrackingNet [37] and LaSOT [13].

5.2 Main Results

VOT2016. The VOT2016 benchmark dataset comprises 60 video sequences. In this dataset, tracking performance is evaluated using three key metrics: Robustness (Rob), Accuracy (Acc), and Expected Average Overlap (EAO) as documented in [29]. As shown in Tab. 1, our proposed Diff-Tracker consistently

Table 1: We provide the evaluation results on the TrackingNet [37], VOT2016 [29], and VOT2018 [30] benchmark datasets. In this table, “Unsup” is an abbreviation for the unsupervised learning.

Tracker	Unsup	TrackingNet			VOT2016			VOT2018		
		Suc \uparrow	Pre \uparrow	NPre \uparrow	EAO \uparrow	Acc \uparrow	Rob \downarrow	EAO \uparrow	Acc \uparrow	Rob \downarrow
SiamFC [1]	No	0.571	0.533	0.663	0.235	0.532	0.461	0.188	0.503	0.585
DaSiamRPN [70]	No	-	-	-	0.411	0.610	0.220	0.326	0.560	0.340
SiamRPN++ [31]	No	0.733	0.694	0.800	-	-	-	0.414	0.600	0.234
ATOM [8]	No	0.703	0.648	0.711	-	-	-	0.401	0.590	0.204
DiMP [2]	No	0.740	0.687	0.801	-	-	-	0.440	0.597	0.153
KCF [20]	Yes	0.447	0.419	0.546	0.192	0.489	0.569	0.135	0.447	0.773
ECO [9]	Yes	0.561	0.489	0.621	0.375	0.550	0.569	0.280	0.270	0.480
S2SiamFC [47]	Yes	-	-	-	0.215	0.493	0.639	0.180	0.463	0.782
LUDT+ [51]	Yes	0.563	0.495	0.633	0.299	0.570	0.331	0.230	0.490	0.412
USOT [68]	Yes	0.599	0.551	0.682	0.351	0.593	0.336	0.290	0.564	0.435
USOT* [68]	Yes	0.616	0.566	0.691	0.402	0.600	0.233	0.344	0.578	0.304
ULAST*-off [46]	Yes	0.649	0.585	0.725	0.397	0.599	0.224	0.347	0.569	0.304
ULAST*-on [46]	Yes	0.654	0.592	0.732	0.417	0.603	0.214	0.355	0.571	0.286
Ours	Yes	0.675	0.614	0.751	0.430	0.605	0.206	0.365	0.580	0.273

improves upon all reported metrics as compared to other unsupervised approaches [46, 68]. Moreover, compared to supervised methods [1, 8], our method still obtains competitive results.

TrackingNet. TrackingNet represents a comprehensive large-scale benchmark tailored for evaluating tracking performance in unconstrained environments, comprising more than 30,000 video sequences. Beyond the precision and success metrics, TrackingNet introduces an additional metric known as normalized precision (NPre). Specifically, it includes a designated test set of 511 videos. As indicated in Tab. 1, the performance of our Diff-Tracker surpasses all other unsupervised trackers.

VOT2018. The testing sequences in VOT2018 present increased challenges in comparison to those in VOT2016. The evaluation metrics used to assess the tracker’s performance on this dataset are the same as those used in VOT2016. Our proposed method is compared against representative supervised and unsupervised trackers, with the evaluation results displayed in Tab. 1. Our Diff-tracker achieves the highest tracking performance among the evaluated unsupervised trackers, according to the three evaluation metrics.

OTB2015. The OTB2015 dataset encompasses 100 video sequences, featuring a diverse array of targets. The evaluation metrics in the OTB2015 benchmark are the success score and precision score. The experimental results are presented in Tab. 2. We surpass prior unsupervised approaches across all metrics, demonstrating the effectiveness of our Diff-Tracker.

Table 2: We provide the evaluation results on the OTB2015 [53] and LaSOT [13] benchmark datasets.

Tracker	Unsup	OTB2015		LaSOT	
		Suc \uparrow	Pre \uparrow	Suc \uparrow	Pre \uparrow
SiamFC [1]	No	0.582	0.771	0.336	0.339
SiamRPN [32]	No	0.637	0.851	0.411	0.380
SiamRPN++ [31]	No	0.696	0.923	0.495	0.493
KCF [20]	Yes	0.485	0.696	0.178	0.166
DSST [11]	Yes	0.518	0.689	0.207	0.189
LUDT+ [51]	Yes	0.639	0.843	0.305	0.288
USOT [68]	Yes	0.589	0.806	0.337	0.323
USOT* [68]	Yes	0.574	0.775	0.358	0.340
ULAST*-off [46]	Yes	0.645	0.878	0.468	0.448
ULAST*-on [46]	Yes	0.648	0.879	0.471	0.451
Ours	Yes	0.661	0.898	0.486	0.472

LaSOT. LaSOT is an extensive annotated collections within the tracking community, comprising 280 extended video sequences. The evaluation metric for this benchmark aligns with that of the OTB2015. We achieve the state-of-the-art performance on the LaSOT dataset, as presented in Tab. 2. For more detailed analysis, please refer to our supplementary.

5.3 Ablation Study

Table 3: We evaluate the impact of the attention harmonization method and online prompt updater, utilizing the VOT2018 [30] benchmark dataset. In the table, “Ours (w/o Attention Harmonization Method)” refers to our method without the designed attention harmonization method.

	EAO \uparrow	Acc \uparrow	Rob \downarrow
Ours (w/o attention harmonization method)	0.359	0.577	0.280
Ours (w/o online prompt updater)	0.349	0.571	0.292
Ours	0.365	0.580	0.273

Attention harmonization. The impact of the harmonization of the self-attention map and the cross-attention map is also explored. Without using the attention harmonization method, we directly compute the MSE loss between the cross-attention map M_c , which is extracted from UNet network of the diffusion model, and the generated GT cross-attention map. As shown in Tab. 3, the addition of the attention harmonization method leads to a performance improvement, showing its efficacy.

Online prompt updater. As shown in Tab. 3, we also verify the impact of the designed online prompt updater in our method. Without the online prompt updater, our Diff-Tracker directly utilizes the learned initial prompt tracking for target tracking in the following frames. From this table, we can observe that the performance of our method decreases in the absence of the online prompt updater, demonstrating that online prompt updater is effective in utilizing the motion information to update the learned prompt, thereby better representing the target.

Table 4: We evaluate the impact of the long-term motion information and short-term motion information, utilizing the VOT2018 [30] benchmark dataset. In the table, “Ours (w/o long-term motion information)” refers to our online prompt updater without the long-term motion information.

	EAO \uparrow	Acc \uparrow	Rob \downarrow
Ours (w/o long-term motion information)	0.355	0.572	0.286
Ours (w/o short-term motion information)	0.360	0.577	0.281
Ours	0.365	0.580	0.273

Long-term and short-term motion information. We explore the impact of long-term and short-term motion information (see Tab. 4). The experimental results show that our method performs best when both long-term and short-term motion information are utilized, which validates the effectiveness of combining these two types of motion information for online updating of the prompt representing the target. For more detailed analysis, please refer to our supplementary.

6 Conclusion

To tackle the challenging task of the unsupervised visual tracking, we develop Diff-Tracker, which leverages the powerful capabilities of the pre-trained text-to-image diffusion models. The Diff-Tracker learns an initial prompt representing the tracking target through a designed initial prompt learner and updates the learned prompt based on the target’s motion using a proposed online prompt updater. Through extensive experiments, we observe that our Diff-Tracker outperforms state-of-the-art unsupervised trackers across five widely used visual tracking benchmarks.

Acknowledgements

This research is supported by the Ministry of Education, Singapore, under the AcRF Tier 2 Projects (MOE-T2EP20222-0009 and MOE-T2EP20123-0014), and the National Research Foundation Singapore through its AI Singapore Programme (AISG-100E-2023-121).

References

1. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. In: *Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8-10 and 15-16, 2016, Proceedings, Part II 14*. pp. 850–865. Springer (2016)
2. Bhat, G., Danelljan, M., Gool, L.V., Timofte, R.: Learning discriminative model prediction for tracking. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 6182–6191 (2019)
3. Budiharto, W., Irwansyah, E., Suroso, J.S., Gunawan, A.A.S.: Design of object tracking for military robot using pid controller and computer vision. *ICIC Express Letters* **14**(3), 289–294 (2020)
4. Chen, J., Ai, Y., Qian, Y., Zhang, W.: A novel siamese attention network for visual object tracking of autonomous vehicles. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering* **235**(10-11), 2764–2775 (2021)
5. Chen, J., Lv, Z., Wu, S., Lin, K.Q., Song, C., Gao, D., Liu, J.W., Gao, Z., Mao, D., Shou, M.Z.: Videollm-online: Online video large language model for streaming video. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18407–18418 (2024)
6. Cheng, X., Xiong, H., Fan, D.P., Zhong, Y., Harandi, M., Drummond, T., Ge, Z.: Implicit motion handling for video camouflaged object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13864–13873 (2022)
7. Cui, Y., Jiang, C., Wang, L., Wu, G.: Mixformer: End-to-end tracking with iterative mixed attention. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 13608–13618 (2022)
8. Danelljan, M., Bhat, G., Khan, F.S., Felsberg, M.: Atom: Accurate tracking by overlap maximization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4660–4669 (2019)
9. Danelljan, M., Bhat, G., Shahbaz Khan, F., Felsberg, M.: Eco: Efficient convolution operators for tracking. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 6638–6646 (2017)
10. Danelljan, M., Gool, L.V., Timofte, R.: Probabilistic regression for visual tracking. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 7183–7192 (2020)
11. Danelljan, M., Häger, G., Khan, F., Felsberg, M.: Accurate scale estimation for robust visual tracking. In: *British machine vision conference, Nottingham, September 1-5, 2014*. Bmva Press (2014)
12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
13. Fan, H., Lin, L., Yang, F., Chu, P., Deng, G., Yu, S., Bai, H., Xu, Y., Liao, C., Ling, H.: Lasot: A high-quality benchmark for large-scale single object tracking. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5374–5383 (2019)
14. Fang, J., Li, Z., Xue, J.: Spatial-sequential-spectral context awareness tracking. In: *2017 IEEE International Conference on Image Processing (ICIP)*. pp. 2582–2586. IEEE (2017)

15. Foo, L.G., Gong, J., Rahmani, H., Liu, J.: Distribution-aligned diffusion for human mesh recovery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9221–9232 (2023)
16. Gao, M., Jin, L., Jiang, Y., Guo, B.: Manifold siamese network: A novel visual tracking convnet for autonomous vehicles. *IEEE Transactions on Intelligent Transportation Systems* **21**(4), 1612–1623 (2019)
17. Gong, J., Foo, L.G., Fan, Z., Ke, Q., Rahmani, H., Liu, J.: Diffpose: Toward more reliable 3d pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13041–13051 (2023)
18. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
19. He, Y., Xu, X., Zhang, J., Shen, F., Yang, Y., Shen, H.T.: Modeling two-stream correspondence for visual sound separation. *IEEE Transactions on Circuits and Systems for Video Technology* **32**(5), 3291–3302 (2021)
20. Henriques, J.F., Caseiro, R., Martins, P., Batista, J.: High-speed tracking with kernelized correlation filters. *IEEE transactions on pattern analysis and machine intelligence* **37**(3), 583–596 (2014)
21. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626* (2022)
22. Hu, H., Gu, J., Zhang, Z., Dai, J., Wei, Y.: Relation networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3588–3597 (2018)
23. Huang, L., Zhao, X., Huang, K.: 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)
24. Hui, X., Wu, Q., Rahmani, H., Liu, J.: Class-agnostic object counting with text-to-image diffusion model. In: European Conference on Computer Vision. Springer (2024)
25. Kawar, B., Zada, S., Lang, O., Tov, O., Chang, H., Dekel, T., Mosseri, I., Irani, M.: Imagic: Text-based real image editing with diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6007–6017 (2023)
26. Khachatryan, L., Movsisyan, A., Tadevosyan, V., Henschel, R., Wang, Z., Navasardyan, S., Shi, H.: Text2video-zero: Text-to-image diffusion models are zero-shot video generators. *arXiv preprint arXiv:2303.13439* (2023)
27. Khani, A., Taghanaki, S.A., Sanghi, A., Amiri, A.M., Hamarneh, G.: Slime: Segment like me. *arXiv preprint arXiv:2309.03179* (2023)
28. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
29. Kristan, M., Leonardis, A., Matas, J., Felsberg, M., Pflugfelder, R., Čehovin, L., et al: The visual object tracking vot2016 challenge results. In: *Computer Vision – ECCV 2016 Workshops*. pp. 777–823. Springer International Publishing (2016)
30. Kristan, M., Leonardis, A., Matas, J., Felsberg, M., Pflugfelder, R., Čehovin Zajc, L., et al: The sixth visual object tracking vot2018 challenge results. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops* (2018)
31. Li, B., Wu, W., Wang, Q., Zhang, F., Xing, J., Yan, J.: Siamrpn++: Evolution of siamese visual tracking with very deep networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4282–4291 (2019)

32. Li, B., Yan, J., Wu, W., Zhu, Z., Hu, X.: High performance visual tracking with siamese region proposal network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8971–8980 (2018)
33. Lin, L., Fan, H., Zhang, Z., Xu, Y., Ling, H.: Swintrack: A simple and strong baseline for transformer tracking. *Advances in Neural Information Processing Systems* **35**, 16743–16754 (2022)
34. Liu, L., Zhang, J., He, R., Liu, Y., Wang, Y., Tai, Y., Luo, D., Wang, C., Li, J., Huang, F.: Learning by analogy: Reliable supervision from transformations for unsupervised optical flow estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6489–6498 (2020)
35. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9298–9309 (2023)
36. Mayer, C., Danelljan, M., Bhat, G., Paul, M., Paudel, D.P., Yu, F., Van Gool, L.: Transforming model prediction for tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8731–8740 (2022)
37. Muller, M., Bibi, A., Giancola, S., Alsubaihi, S., Ghanem, B.: Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In: Proceedings of the European conference on computer vision (ECCV). pp. 300–317 (2018)
38. Papanikolopoulos, N.P., Khosla, P.K., Kanade, T.: Visual tracking of a moving target by a camera mounted on a robot: A combination of control and vision. *IEEE transactions on robotics and automation* **9**(1), 14–35 (1993)
39. Paul, M., Danelljan, M., Mayer, C., Van Gool, L.: Robust visual tracking by segmentation. In: European Conference on Computer Vision. pp. 571–588. Springer (2022)
40. Peng, D., Ke, Q., Lei, Y., Liu, J.: Unsupervised domain adaptation via domain-adaptive diffusion. *arXiv preprint arXiv:2308.13893* (2023)
41. Peng, D., Zhang, Z., Hu, P., Ke, Q., Yau, D., Liu, J.: Harnessing text-to-image diffusion models for category-agnostic pose estimation. In: European Conference on Computer Vision. Springer (2024)
42. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
43. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
44. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115**, 211–252 (2015)
45. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems* **35**, 36479–36494 (2022)
46. Shen, Q., Qiao, L., Guo, J., Li, P., Li, X., Li, B., Feng, W., Gan, W., Wu, W., Ouyang, W.: Unsupervised learning of accurate siamese tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8101–8110 (2022)
47. Sio, C.H., Ma, Y.J., Shuai, H.H., Chen, J.C., Cheng, W.H.: S2siamfc: Self-supervised fully convolutional siamese network for visual tracking. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 1948–1957 (2020)

48. Tan, M., Pang, R., Le, Q.V.: Efficientdet: Scalable and efficient object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10781–10790 (2020)
49. Tang, L., Jia, M., Wang, Q., Phoo, C.P., Hariharan, B.: Emergent correspondence from image diffusion. arXiv preprint arXiv:2306.03881 (2023)
50. Wang, N., Song, Y., Ma, C., Zhou, W., Liu, W., Li, H.: Unsupervised deep tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1308–1317 (2019)
51. Wang, N., Zhou, W., Song, Y., Ma, C., Liu, W., Li, H.: Unsupervised deep representation learning for real-time tracking. *International Journal of Computer Vision* **129**, 400–418 (2021)
52. Wong, B., Chen, J., Wu, Y., Lei, S.W., Mao, D., Gao, D., Shou, M.Z.: Assistq: Affordance-centric question-driven task completion for egocentric assistant. In: European Conference on Computer Vision. pp. 485–501. Springer (2022)
53. Wu, Y., Lim, J., Yang, M.H.: Object tracking benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **37**(9), 1834–1848 (2015)
54. Xie, S., Zhang, Z., Lin, Z., Hinz, T., Zhang, K.: Smartbrush: Text and shape guided object inpainting with diffusion model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22428–22437 (2023)
55. Xie, X., Cheng, G., Wang, J., Yao, X., Han, J.: Oriented r-cnn for object detection. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3520–3529 (2021)
56. Xu, L., Huang, H., Liu, J.: Sutd-trafficqa: A question answering benchmark and an efficient network for video reasoning over traffic events. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9878–9888 (2021)
57. Xu, L., Huang, M.H., Shang, X., Yuan, Z., Sun, Y., Liu, J.: Meta compositional referring expression segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19478–19487 (2023)
58. Xu, N., Yang, L., Fan, Y., Yang, J., Yue, D., Liang, Y., Price, B., Cohen, S., Huang, T.: Youtube-vos: Sequence-to-sequence video object segmentation. In: Proceedings of the European conference on computer vision (ECCV). pp. 585–601 (2018)
59. Yan, B., Peng, H., Fu, J., Wang, D., Lu, H.: Learning spatio-temporal transformer for visual tracking. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10448–10457 (2021)
60. Ye, B., Chang, H., Ma, B., Shan, S., Chen, X.: Joint feature learning and relation modeling for tracking: A one-stream framework. In: European Conference on Computer Vision. pp. 341–357. Springer (2022)
61. Yu, C., Wang, J., Peng, C., Gao, C., Yu, G., Sang, N.: Bisenet: Bilateral segmentation network for real-time semantic segmentation. In: Proceedings of the European conference on computer vision (ECCV). pp. 325–341 (2018)
62. Yu, Y., Xiong, Y., Huang, W., Scott, M.R.: Deformable siamese attention networks for visual object tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6728–6737 (2020)
63. Yuan, D., Chang, X., Huang, P.Y., Liu, Q., He, Z.: Self-supervised deep correlation tracking. *IEEE Transactions on Image Processing* **30**, 976–985 (2020)
64. Zhang, Z., Zhou, C., Tu, Z.: Distilling inter-class distance for semantic segmentation. arXiv preprint arXiv:2205.03650 (2022)
65. Zhang, Z., Zhou, Y., Gong, J., Liu, J., Tu, Z.: Instance temperature knowledge distillation. arXiv preprint arXiv:2407.00115 (2024)

66. Zhao, Q., Dai, Y., Li, H., Hu, W., Zhang, F., Liu, J.: Ltgc: Long-tail recognition via leveraging llms-driven generated content. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19510–19520 (2024)
67. Zhao, Q., Huang, Y., Hu, W., Zhang, F., Liu, J.: Mixpro: Data augmentation with maskmix and progressive attention labeling for vision transformer. arXiv preprint arXiv:2304.12043 (2023)
68. Zheng, J., Ma, C., Peng, H., Yang, X.: Learning to track objects from unlabeled videos. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 13546–13555 (2021)
69. Zhu, G., Wang, J., Zhao, C., Lu, H.: Weighted part context learning for visual tracking. *IEEE Transactions on Image Processing* **24**(12), 5140–5151 (2015)
70. Zhu, Z., Wang, Q., Li, B., Wu, W., Yan, J., Hu, W.: Distractor-aware siamese networks for visual object tracking. In: Proceedings of the European conference on computer vision (ECCV). pp. 101–117 (2018)