

Exploring Enhanced Contextual Information for Video-Level Object Tracking

Ben Kang^{1,2}, Xin Chen¹, Simiao Lai¹, Yang Liu^{1*}, Yi Liu³, Dong Wang^{1,2}

¹ Dalian University of Technology

² Ningbo Institute of Dalian University of Technology

³ Baidu Inc.

kangben@mail.dlut.edu.cn, chenxin3131@mail.dlut.edu.cn, laisimiao@mail.dlut.edu.cn, ly@dlut.edu.cn, liuyi22@baidu.com, wdice@dlut.edu.cn

Abstract

Contextual information at the video level has become increasingly crucial for visual object tracking. However, existing methods typically use only a few tokens to convey this information, which can lead to information loss and limit their ability to fully capture the context. To address this issue, we propose a new video-level visual object tracking framework called **MCITrack**. It leverages Mamba’s hidden states to continuously record and transmit extensive contextual information throughout the video stream, resulting in more robust object tracking. The core component of MCITrack is the Contextual Information Fusion module, which consists of the mamba layer and the cross-attention layer. The mamba layer stores historical contextual information, while the cross-attention layer integrates this information into the current visual features of each backbone block. This module enhances the model’s ability to capture and utilize contextual information at multiple levels through deep integration with the backbone. Experiments demonstrate that MCITrack achieves competitive performance across numerous benchmarks. For instance, it gets 76.6% AUC on LaSOT and 80.0% AO on GOT-10k, establishing a new state-of-the-art performance.

Code — <https://github.com/kangben258/MCITrack>

Introduction

Visual object tracking is a crucial task in computer vision, aiming to identify an object in the initial frame of a video and predict its location in subsequent frames. Existing tracking methods (Danelljan et al. 2019; Chen et al. 2020, 2022; Zhu et al. 2023; Xie et al. 2022) typically employ image-level matching techniques, using template and search region images for feature matching and then making predictions based on the matching results. These methods primarily focus on the initial appearance of the target and do not utilize the contextual information available in the video sequence.

To improve performance, some methods (Yan et al. 2021a; Mayer et al. 2022; Chen et al. 2023b) use dynamic templates to record the appearance information of objects. As shown in Figure 1 (a), these methods obtain more appearance information and enhance the model’s robustness by

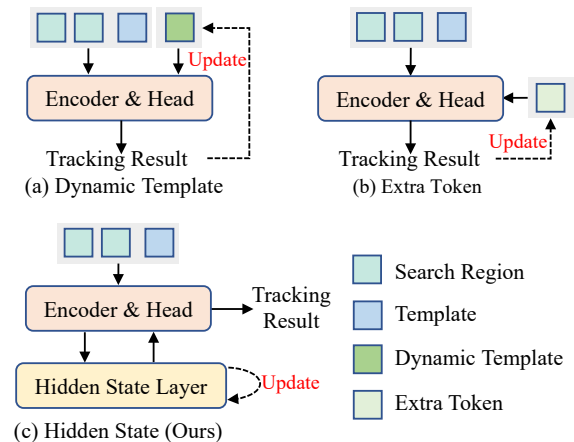


Figure 1: Comparison of different contextual information propagation methods. (a) Using dynamic templates. (b) Using extra tokens. (c) Using hidden state layers.

updating the dynamic template. While these methods have achieved some success, they still primarily focus on the appearance information of the target and do not fully utilize the contextual information in the video sequence.

Recently, some works (Chang et al. 2024; Zheng et al. 2024; Bai et al. 2024; Shi et al. 2024; Xie et al. 2024) have begun to explore video-level object tracking. Unlike previous methods focusing solely on the target’s appearance, these works utilize more contextual information to enhance model performance. As shown in Figure 1 (b), these methods typically use a small number of extra tokens to capture and transmit contextual information, continuously updating as the video sequence progresses. Although these methods show superior performance compared to earlier approaches, the limited number of tokens can only transmit a limited amount of contextual information, leading to information loss. Therefore, efficiently transmitting richer contextual information remains an urgent challenge.

Temporal models like LSTM (Graves 2012) and state space models (Gu et al. 2021) are adept at preserving crucial information in their hidden states, making them suitable for recording important contextual information in video se-

*Corresponding author.

quences. To address the issues mentioned above, we propose a new framework for transmitting contextual information in video-level object tracking. As shown in Figure 1 (c), we use hidden state layers to store and transmit richer contextual information. Specifically, we introduce MCITrack, which leverages hidden states from Mamba (Gu and Dao 2023) to record and transmit contextual information. Compared to using a small number of additional tokens, MCITrack can record and transmit more implicit contextual information, thereby enhancing the model’s performance.

The core module of MCITrack is the Contextual Information Fusion (CIF) module, which consists of mamba and cross-attention layers. The CIF module inputs multi-level features from the backbone blocks of the current frame into the mamba layers, where hidden states record crucial information. As the next frame is processed, the cross-attention layers integrate this recorded information into each backbone block. This deep integration allows the CIF module to effectively extract and utilize contextual information at multiple levels, aiding the model in making accurate predictions. Throughout the video, MCITrack continuously updates the hidden states, ensuring the effective transmission of contextual information as the sequence progresses.

Extensive experiments demonstrate the effectiveness of our contextual information transmission framework. MCITrack achieves state-of-the-art performance on multiple datasets. For example, compared to the recent state-of-the-art model ODTrack-B384 (Zheng et al. 2024), MCITrack-B224 improves the AUC on LaSOT (Fan et al. 2019) by 2.1% (75.3% *vs.* 73.2%) with lower computational cost (38G *vs.* 73G), fewer parameters (88M *vs.* 92M), and a lower resolution (256 *vs.* 384). Even when compared to the larger ODTrack-L384, MCITrack-B224 achieves a 1.3% higher AUC on LaSOT. Notably, our largest model, MCITrack-L384, reaches an unprecedented AUC of 76.6% on LaSOT. These results highlight the effectiveness of our framework in transmitting critical contextual information, enhancing prediction accuracy beyond previous methods. The contributions of this paper are summarized as follows:

- We introduce a new method for transmitting contextual information in video-level object tracking. Compared with previous methods, our approach effectively transmits richer and more crucial contextual information.
- Based on the proposed framework, we develop a new family of video-level object tracking models named MCITrack. Experiments demonstrate the effectiveness of this framework, with MCITrack achieving state-of-the-art performance across multiple datasets.

Related Work

Visual Tracking. Current trackers generally fall into two categories: image-level trackers that rely solely on appearance information and video-level trackers that incorporate more contextual information. Most trackers (Tao, Gavves, and Smeulders 2016; Bertinetto et al. 2016; Li et al. 2018, 2019; Zhang and Peng 2019; Xu et al. 2020; Voigtlaender et al. 2020; Gao et al. 2022; Guo et al. 2022) belong to the first category. These trackers use only the initial appearance

information of the target to match with the current frame and then predict the target’s location based on the matching results through a tracking head. Such trackers often struggle with challenges like object deformation and blurriness. Some trackers (Fu et al. 2021; Wang et al. 2021a; Yan et al. 2021a; Mayer et al. 2021; Chen et al. 2023a) introduce additional templates, continuously updating them to allow the model to extract more appearance information about the target to address these challenges. Despite these trackers having achieved some success, they still rely solely on appearance for image-level matching to locate the target. They do not utilize the contextual information in the video.

Recently, many researchers have begun exploring video-level object tracking to better utilize contextual information for assisting tracking. TCTrack (Cao et al. 2022) aggregates contextual information through online temporal adaptive convolution and temporal adaptive feature map refinement. VideoTrack (Xie et al. 2023) employs a video transformer backbone to integrate contextual information. Besides these methods, scholars have also started exploring the use of additional tokens for this purpose. ODTrack (Zheng et al. 2024) employs a token sequence propagation paradigm to densely associate contextual relationships across video frames. ARTrack (Wei et al. 2023) and ARTrackV2 (Bai et al. 2024) encode previously predicted object coordinate sequences into tokens to assist the current prediction. AQA-Track (Xie et al. 2024) uses learnable autoregressive queries to capture contextual information. EVPTrack (Shi et al. 2024) utilizes spatio-temporal markers to propagate information across consecutive frames.

Compared with previous methods, video-level object tracking models incorporating contextual information have significantly improved tracking performance. However, these models typically rely on a limited number of additional tokens to capture contextual information, which restricts the ability to capture comprehensive contextual information, leading to information loss. In contrast, our MCITrack efficiently captures richer and more critical contextual information, enhancing the model’s performance.

State Space Model. The State Space Model (SSM) is proposed for handling sequence tasks in NLP (Gu et al. 2021; Gu, Goel, and Ré 2022) and has shown potential in addressing long-range dependency problems. Mamba (Gu and Dao 2023) addresses SSM’s inability to perform content-based reasoning by making SSM parameters functions of the input. Mamba has demonstrated the powerful capability of SSM in handling long sequences with higher throughput compared to Transformers. Subsequently, it has also become popular in the computer vision field. In image classification, ViM (Zhu et al. 2024) and VMamba (Liu et al. 2024) have achieved performance comparable to ViT (Dosovitskiy et al. 2020) with fewer parameters. In the medical image domain, SSM has found widespread applications. For instance, MambaUNet (Wang et al. 2024) incorporates the Mamba structure into the UNet (Ronneberger, Fischer, and Brox 2015) model for medical image segmentation. In this work, we explore how to efficiently capture and transmit contextual information in video sequences using the hidden states in Mamba.

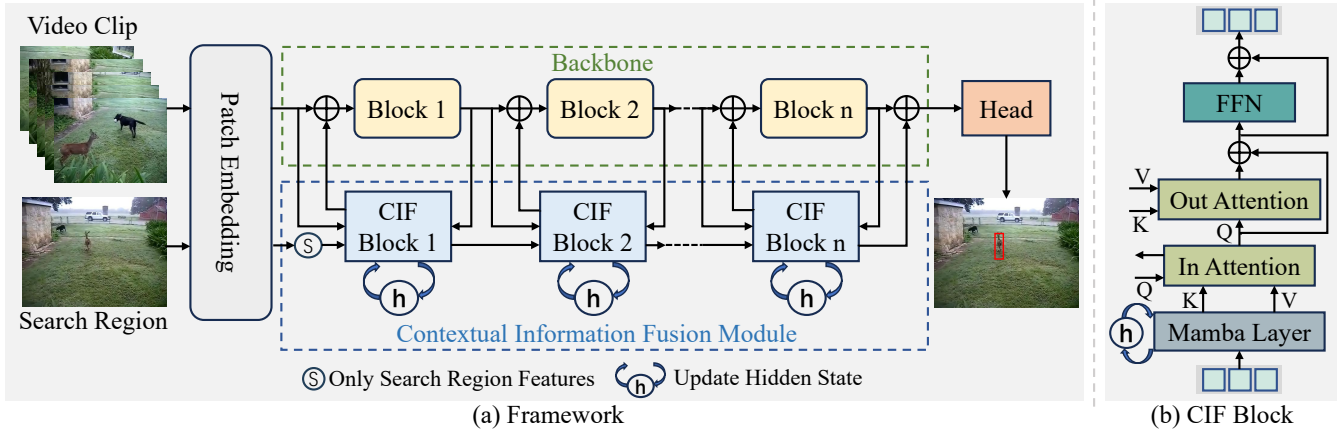


Figure 2: (a) Framework of the proposed MCITrack. The key components include the backbone for visual feature extraction, the contextual information fusion module for recording and transmitting contextual information, and the prediction head. (b) Architecture of the proposed Contextual Information Fusion(CIF) block.

MCITrack

This section presents MCITrack in detail. First, we provide an overview of MCITrack. Next, we describe the model architecture, including the backbone and the Contextual Information Fusion (CIF) module. Finally, we introduce the prediction head and training objective function.

Overview

The overall architecture of MCITrack is shown in Figure 2 (a). It consists of three parts: a backbone for visual feature extraction, a Contextual Information Fusion (CIF) module for storing and transmitting contextual information, and a prediction head for making predictions. MCITrack takes a video clip and a search region as input. First, the video clip and search region are divided into patches through patch embedding, and then these patches are concatenated along the spatial dimension and fed into the backbone for feature extraction. The backbone is composed of N blocks, with each block paired with a corresponding CIF block. The CIF blocks integrate the historical contextual information into their associated backbone blocks, enhancing the accuracy of visual feature extraction based on the historical contextual information. Simultaneously, the CIF blocks update the hidden states based on the current backbone output. Finally, the backbone, guided by contextual information, outputs more precise visual features, which are then passed to the prediction head to obtain the tracking results.

Model Architecture

Backbone. We use Fast-iTPN (Tian et al. 2024) as our backbone. Similar to ViT (Dosovitskiy et al. 2020), Fast-iTPN consists of Transformer layers. However, compared to ViT, Fast-iTPN is "narrower and deeper", featuring a smaller hidden dimension and more network layers. This deeper structure enables the CIF module to capture a broader range of information, including both shallow-level details and deeper semantic information. Our model takes a video clip

$\mathbf{V} \in \mathbb{R}^{N \times 3 \times H_z \times W_z}$ and a search region $\mathbf{X} \in \mathbb{R}^{3 \times H_x \times W_x}$ as inputs. First, both \mathbf{V} and \mathbf{X} are downsampled using convolutional layers with a stride of 4. The downsampled \mathbf{V} and \mathbf{X} are processed through MLP layers and two convolutional merging layers. This process segments \mathbf{V} and \mathbf{X} into patches, resulting in $\mathbf{F}_v \in \mathbb{R}^{N \times C \times \frac{H_z}{16} \times \frac{W_z}{16}}$ and $\mathbf{F}_x \in \mathbb{R}^{C \times \frac{H_x}{16} \times \frac{W_x}{16}}$. These patches are expanded and concatenated along the spatial dimension to form $\mathbf{F}_{vx} \in \mathbb{R}^{C \times L}$ ($L = N \times \frac{H_z}{16} \times \frac{W_z}{16} + \frac{H_x}{16} \times \frac{W_x}{16}$), which are then input into the Transformer layers for visual feature extraction. To enable the CIF module to capture more information at multiple levels, we divide the backbone's Transformer layers into N (N is 4 by default) blocks, each paired with a corresponding CIF block. Before the features are input into each block, the historical contextual information from the CIF block is integrated, enhancing the visual feature extraction process. This integration of contextual information aids the backbone in extracting more accurate features for subsequent predictions by the prediction head.

Contextual Information Fusion Module. The Contextual Information Fusion (CIF) module is the core component of MCITrack. It stores and integrates contextual information into the backbone to enhance visual feature extraction. The CIF module consists of N CIF blocks, each corresponding to a backbone block. Each CIF block integrates contextual information into the backbone features before the corresponding backbone block performs feature extraction, and updates the hidden states based on the current frame's features. This deep integration enables the CIF module to extract information at multiple levels, providing richer implicit contextual information to improve the model's predictions. The structure of the CIF block is shown in Figure 2 (b) and primarily includes a mamba layer and two cross-attention layers. The mamba layer is responsible for storing contextual information. The in-attention layer integrates the contextual information from the CIF block into the backbone, while the out-attention layer extracts current frame information from

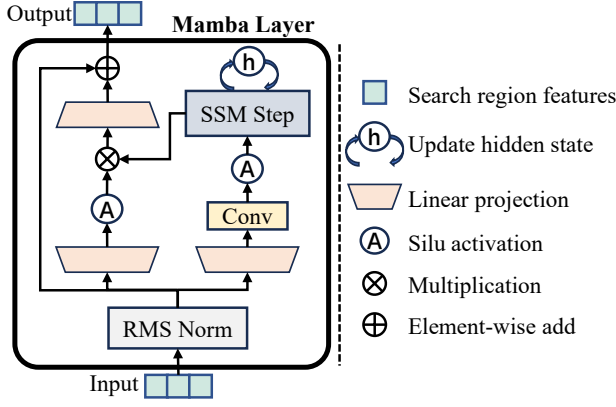


Figure 3: The detailed mamba layer in the CIF block.

the backbone’s features into the CIF block. Specifically, the output from the previous CIF block, \mathbf{F}_c^{i-1} , is used as input for the current CIF block. \mathbf{F}_c^{i-1} contains information extracted from the features of the previous backbone block’s output, \mathbf{F}_{vx}^{i-1} . First, \mathbf{F}_c^{i-1} passes through the mamba layer, which extracts the historical contextual information, resulting in \mathbf{F}_h^i . Simultaneously, the mamba layer updates its hidden states based on \mathbf{F}_c^{i-1} , thereby updating the contextual information. Next, \mathbf{F}_{vx}^{i-1} is used as Q , while \mathbf{F}_h^i serves as K and V in the in-attention layer to obtain $\mathbf{F}_{vx}^{i-1'}$, which incorporates the previous contextual information. Then, $\mathbf{F}_{vx}^{i-1'}$ is input into the backbone for further feature extraction, resulting in \mathbf{F}_{vx}^i . Subsequently, \mathbf{F}_{vx}^i is used as K and V , while \mathbf{F}_h^i is used as Q in the residual connected out-attention and FFN to extract the current frame’s information into the CIF block, resulting in the current CIF block’s output, \mathbf{F}_c^i . The processing of the CIF module can be summarized as follows:

$$\begin{aligned}
 \mathbf{F}_h^i, \mathbf{H}_t^i &= \text{Mamba}_i(\mathbf{F}_c^{i-1}, \mathbf{H}_{t-1}^i), \\
 \mathbf{F}_{vx}^i &= \text{Block}_i(\mathbf{F}_{vx}^{i-1} + \text{Att}_{\text{in}}^i(\mathbf{F}_{vx}^{i-1}, \mathbf{F}_h^i, \mathbf{F}_h^i)), \\
 \mathbf{F}_c^{i'} &= \text{Att}_{\text{out}}^i(\mathbf{F}_h^i, \mathbf{F}_{vx}^i, \mathbf{F}_{vx}^i) + \mathbf{F}_h^i, \\
 \mathbf{F}_c^i &= \mathbf{F}_c^{i'} + \text{FFN}_i(\mathbf{F}_c^{i'}),
 \end{aligned} \tag{1}$$

Where, Mamba_i , Att_{in}^i , $\text{Att}_{\text{out}}^i$ and FFN_i represent the mamba layer, in-attention, out-attention, and FFN layer, respectively, in the i_{th} CIF block, Block_i represents the i_{th} backbone block, \mathbf{H}_t^i and \mathbf{H}_{t-1}^i represent the hidden states of the i_{th} CIF block at time t and $t - 1$, respectively.

Mamba Layer. The mamba layer is a key component of the CIF block, responsible for storing contextual information. The detailed structure of the mamba layer is shown in Figure 3. First, the input \mathbf{F}_i undergoes RMS Norm (Zhang and Sennrich 2019) for normalization, resulting in \mathbf{F}_{in} . Next, \mathbf{F}_{in} is processed through two branches: x and z . In the x branch, \mathbf{F}_{in} is expanded in dimension via a linear projection and then passed through a convolutional layer to extract local features. After activation by the SiLU function (Elfwing, Uchibe, and Doya 2018), it is fed into the SSM to extract and update the stored contextual information. The output of the x branch is denoted as \mathbf{F}_x . In the z branch, \mathbf{F}_{in} is dimen-

sionally expanded and activated by the SiLU function. It is then weighted by \mathbf{F}_x and reduced in dimension to produce the output \mathbf{F}'_i of the z branch. Finally, \mathbf{F}'_i is combined with \mathbf{F}_{in} through a skip connection to produce the output \mathbf{F}_o of the Mamba Layer. This output \mathbf{F}_o integrates the previously stored contextual information.

State Space Model. The State Space Model (SSM) is a critical part of the mamba layer. Inspired by continuous systems, SSM maps the one-dimensional function $x(t) \in \mathbb{R}$ to $y(t) \in \mathbb{R}$ through the system’s hidden state function $h(t) \in \mathbb{R}^N$ (N represents the state size). SSM can be represented by a set of linear ordinary differential equations, specifically the state equation and the output equation:

$$\begin{aligned}
 h'(t) &= \mathbf{A}h(t) + \mathbf{B}x(t), \\
 y(t) &= \mathbf{C}h(t) + \mathbf{D}x(t),
 \end{aligned} \tag{2}$$

Where $\mathbf{A} \in \mathbb{R}^{N \times N}$, $\mathbf{B} \in \mathbb{R}^{N \times 1}$, $\mathbf{C} \in \mathbb{R}^{1 \times N}$, and $\mathbf{D} \in \mathbb{R}^1$ are the weighting parameters of the continuous-time SSM. To incorporate a continuous-time SSM in deep learning, it must be discretized. The zero-order hold method (Gu and Dao 2023) is typically used to discretize it, converting continuous parameters into discrete parameters, as follows:

$$\begin{aligned}
 \bar{\mathbf{A}} &= \exp(\Delta \mathbf{A}), \\
 \bar{\mathbf{B}} &= (\Delta \mathbf{A})^{-1}(\exp(\Delta \mathbf{A}) - \mathbf{I}) \cdot \Delta \mathbf{B} \approx \Delta \mathbf{B},
 \end{aligned} \tag{3}$$

Where $\bar{\mathbf{A}}$ and $\bar{\mathbf{B}}$ are the discretized parameters, and Δ is the time scale parameter. After discretization, the SSM used in the mamba layer is represented by the following formula:

$$\begin{aligned}
 h_t &= \bar{\mathbf{A}}h_{t-1} + \bar{\mathbf{B}}x_t, \\
 y_t &= \mathbf{C}h_t + \mathbf{D}x_t,
 \end{aligned} \tag{4}$$

Here, h_{t-1} represents the previous hidden states that store crucial contextual information, h_t is the updated hidden states based on the current input x_t , and y_t is the output that integrates the contextual information. In this way, SSM enables storing and transmitting contextual information.

Head and Loss Function

We use the same classification and regression heads as OS-Track (Ye et al. 2022) for our prediction head. The head consists of three sub-networks, each composed of convolutional layers. These sub-networks produce the following outputs: the classification score $S \in \mathbb{R}^{1 \times \frac{H_x}{16} \times \frac{W_x}{16}}$, bounding box size $B \in \mathbb{R}^{2 \times \frac{H_x}{16} \times \frac{W_x}{16}}$, and offset size $O \in \mathbb{R}^{2 \times \frac{H_x}{16} \times \frac{W_x}{16}}$. Using S , B , and O , we predict the final tracking results. Our objective function includes classification loss, focal loss (Lin et al. 2017), and regression loss, which consists of L1 loss and GloU loss (Rezatofighi et al. 2019). In our video-level training approach, each frame contributes to an individual loss. The total loss can be summarized as follows:

$$\mathcal{L} = \sum_{i=1}^n (\lambda_c \mathcal{L}_{cls}^i + \lambda_l \mathcal{L}_l^i + \lambda_g \mathcal{L}_g^i), \tag{5}$$

Where, \mathcal{L}_{cls}^i , \mathcal{L}_l^i and \mathcal{L}_g^i represent the classification loss, L1 loss, and GloU loss for the i_{th} frame, respectively. \mathcal{L} is the total loss. λ_c , λ_l and λ_g are hyperparameters, with default values $\lambda_c = 1$, $\lambda_l = 5$ and $\lambda_g = 2$.

Method	LaSOT			LaSOT _{ext}			TrackingNet			GOT-10k*		
	AUC	P _{Norm}	P	AUC	P _{Norm}	P	AUC	P _{Norm}	P	AO	SR _{0.5}	SR _{0.75}
MCITrack-B224	75.3	85.6	83.3	54.6	65.7	62.1	86.3	90.9	86.1	77.9	88.2	76.8
MCITrack-S224	<u>73.8</u>	<u>84.2</u>	<u>81.7</u>	<u>52.6</u>	63.6	59.7	<u>85.6</u>	<u>90.2</u>	<u>85.2</u>	76.9	87.0	<u>76.1</u>
MCITrack-T224	71.7	81.5	78.2	51.6	62.7	58.4	84.8	89.4	83.7	74.0	83.9	72.1
ODTrack-B384 (Zheng et al. 2024)	73.2	83.2	80.6	52.4	<u>63.9</u>	<u>60.1</u>	85.1	90.1	84.9	<u>77.0</u>	<u>87.9</u>	75.1
ARTrackV2-256 (Bai et al. 2024)	71.6	80.2	77.2	50.8	61.9	57.7	84.9	89.3	84.5	75.9	85.4	72.7
AQATrack-256 (Xie et al. 2024)	71.4	81.9	78.6	51.2	62.2	58.9	83.8	88.6	83.1	73.8	83.2	72.1
EVPTrack-224 (Shi et al. 2024)	70.4	80.9	77.2	48.7	59.5	55.1	83.5	88.3	-	73.3	83.6	70.7
LoRAT-B224 (Lin et al. 2024)	71.7	80.9	77.3	50.3	61.6	57.1	83.5	87.9	82.1	72.1	81.8	70.7
ARTrack-256 (Wei et al. 2023)	70.4	79.5	76.6	46.4	56.5	52.3	84.2	88.7	83.5	73.5	82.2	70.9
SeqTrack-B256 (Chen et al. 2023a)	69.9	79.7	76.3	49.5	60.8	56.3	83.3	88.3	82.2	74.7	84.7	71.8
VideoTrack-256 (Xie et al. 2023)	70.2	-	76.4	-	-	-	83.8	88.7	83.1	72.9	81.9	69.8
ROMTrack-384 (Cai et al. 2023)	71.4	81.4	78.2	51.3	62.4	58.6	84.1	89.0	83.7	74.2	84.3	72.4
CiteTracker-384 (Li et al. 2023)	69.7	78.6	75.7	-	-	-	84.5	89.0	84.2	74.7	84.3	73.0
OSTrack-256 (Ye et al. 2022)	69.1	78.7	75.2	47.4	57.3	53.3	83.1	87.8	82.0	71.0	80.4	68.2
Mixformer-22k (Cui et al. 2022)	69.2	78.7	74.7	-	-	-	83.1	88.1	81.6	70.7	80.0	67.8
SwinTrack-384 (Lin et al. 2022)	71.3	-	76.5	49.1	-	55.6	84.0	-	82.8	72.4	-	67.8
Comparison of Large-Scale Models												
MCITrack-L384	76.6	86.1	85.0	<u>55.7</u>	<u>66.5</u>	<u>62.9</u>	87.9	92.1	89.2	80.0	<u>88.5</u>	80.2
MCITrack-L224	<u>76.1</u>	86.1	<u>84.1</u>	54.8	65.6	61.6	<u>86.9</u>	<u>91.3</u>	<u>87.4</u>	79.3	89.3	78.7
ODTrack-L384 (Zheng et al. 2024)	74.0	<u>84.2</u>	82.3	53.9	65.4	61.7	86.1	91.0	86.7	78.2	87.2	77.3
ARTrackV2-L384 (Bai et al. 2024)	73.6	82.8	81.1	53.4	63.7	60.2	86.1	90.4	86.2	<u>79.5</u>	87.8	<u>79.6</u>
LoRAT-L378 (Lin et al. 2024)	75.1	84.1	82.0	56.6	69.0	65.1	85.6	89.7	85.4	<u>77.5</u>	86.2	78.1
MixViT-L384 (Cui et al. 2024)	72.4	82.2	80.1	-	-	-	85.4	90.2	85.7	75.7	85.3	75.1
ARTrack-L384 (Wei et al. 2023)	73.1	82.2	80.3	52.8	62.9	59.7	85.6	89.6	86.0	78.5	87.4	77.8
SeqTrack-L384 (Chen et al. 2023a)	72.5	81.5	79.3	50.7	61.6	57.5	85.5	89.8	85.8	74.8	81.9	72.2
TATrack-L384 (He et al. 2023)	71.1	79.1	76.1	-	-	-	85.0	89.3	84.5	-	-	-
CTTrack-L320 (Song et al. 2023)	69.8	78.7	76.2	-	-	-	84.9	89.1	83.5	72.8	81.3	71.5

Table 1: State-of-the-art comparisons on four large-scale benchmarks. We add a symbol * over GOT-10k to indicate that the corresponding models are only trained with the GOT-10k training set. Otherwise, the models are trained with all the training data presented in Sec. Implementation Details. The top two results are highlighted with **bold** and underlined fonts, respectively.

Experiments

Implementation Details

Model. We develop five variants of the MCITrack model with different backbones and input resolutions, as detailed in Table 2. Additionally, we report the model parameters, FLOPs, and inference speed in Table 2. The speed is measured on an Intel Core i7-8700K CPU @3.70GHz with 47GB RAM and a single 2080 Ti GPU. All the models are implemented with Python 3.11 and PyTorch 2.1.2.

Training. Our training data includes LaSOT (Fan et al. 2019), GOT-10k (Huang, Zhao, and Huang 2019), TrackingNet (Muller et al. 2018), COCO (Lin et al. 2014), and VastTrack (Peng et al. 2024). We use a 5-frame video clip and two search regions as inputs. To simulate contextual information transmission, we first input the video clip and one search region, calculate the loss, and store the contextual information in hidden states. Then, we input the second search region and the previous frame’s hidden states, calculate the loss, and sum the losses for backpropagation. We use the AdamW (Loshchilov and Hutter 2018) optimizer with an initial learning rate of 4×10^{-5} for the backbone and 4×10^{-4} for the rest. The weight decay is set to 1×10^{-4} . The model is trained for a total of 300 epochs with 60k sam-

Model	Backbone	Input Resolution	Params (M)	FLOPs (G)	Speed (fps)
MCITrack-T224	Fast-iTPN-T	224×224	32	13	51
MCITrack-S224	Fast-iTPN-S	224×224	45	19	40
MCITrack-B224	Fast-iTPN-B	224×224	88	38	35
MCITrack-L224	Fast-iTPN-L	224×224	287	123	16
MCITrack-L384	Fast-iTPN-L	384×384	287	370	5

Table 2: Details of MCITrack model variants.

ples per epoch, and the learning rate decreases by a factor of 10 after 240 epochs. Training is performed on two 80GB Tesla A800 GPUs with a total batch size of 128.

Inference. During inference, we input a video clip of length 5, consistent with our training process. We maintain a memory bank to store reliable frames. When the update interval T is reached, we use the frames stored in the memory bank to update the video clip. We observe that updating the hidden states when tracking results are inaccurate can introduce erroneous contextual information, which may degrade the model’s performance in subsequent frames. To mitigate this issue, we set a threshold a for updating the hid-

Method	TNL2K	NFS	UAV123
MCITrack-L384	65.3	70.6	<u>71.5</u>
MCITrack-L224	<u>64.3</u>	71.1	70.8
MCITrack-B224	62.9	<u>70.6</u>	70.5
MCITrack-S224	61.9	<u>70.6</u>	69.3
MCITrack-T224	59.4	<u>70.0</u>	69.9
ODTrack-L384 (Zheng et al. 2024)	61.7	-	-
ARTrackV2-256 (Bai et al. 2024)	59.2	67.6	69.9
LoRAT-L378 (Lin et al. 2024)	62.3	66.7	72.5
ARTrack-L384 (Wei et al. 2023)	60.3	67.9	71.2
SeqTrack-L384 (Chen et al. 2023a)	57.8	66.2	68.5
OTrack (Ye et al. 2022)	55.9	66.5	70.7
STARK (Yan et al. 2021a)	-	66.2	68.2
TransT (Chen et al. 2021)	50.7	65.7	69.1
DiMP (Bhat et al. 2019)	44.7	61.8	64.3
Ocean (Zhang et al. 2020)	38.4	49.4	57.4

Table 3: Comparison with state-of-the-art methods on additional benchmarks in AUC score.

den states: the hidden states are only updated if the classification score exceeds the threshold. This approach helps prevent the model from being misled by incorrect information.

State-of-the-Art Comparisons

We compare our MCITrack with SOTA trackers across eight tracking benchmarks, as detailed in Tables 1 and 3.

LaSOT. LaSOT (Fan et al. 2019) is a large-scale, long-term dataset with 1120 training videos and 208 test videos. As shown in Table 1, MCITrack models achieve state-of-the-art results. Specifically, MCITrack-B224 and MCITrack-S224 attain the top two AUC scores of 75.3% and 73.8%, respectively, outperforming the third-place ODTrack-B384 by 2.1% and 0.6%. Furthermore, among the large-scale models, MCITrack-L384 and MCITrack-L224 secure the first and second positions with AUC scores of 76.6% and 76.1%.

LaSOT_{ext}. LaSOT_{ext} (Fan et al. 2021) is an extension of LaSOT, containing 150 videos. As reported in Table 1, our models achieve the top two results. MCITrack-B224 achieves an AUC score of 54.6%, which is 2.2% higher than the previous best, ODTrack-B384, representing a significant improvement. With a large-scale backbone, MCITrack-L384 achieves a competitive AUC score of 55.7%.

TrackingNet. TrackingNet (Muller et al. 2018) is a comprehensive dataset that covers diverse object categories and scenes. As presented in Table 1, MCITrack-B224 and MCITrack-S224 achieve the top two positions with AUC scores of 86.3% and 85.6%, respectively. Additionally, MCITrack-L384 and MCITrack-L224 also secure the top two results among large-scale trackers, with AUC scores of 1.8% and 0.8% higher than the third-place tracker, demonstrating a significant performance advantage.

GOT-10k. GOT-10k (Huang, Zhao, and Huang 2019) test set contains 180 videos covering various common tracking challenges. In line with official guidelines, we only use the GOT-10k training set for model training. As

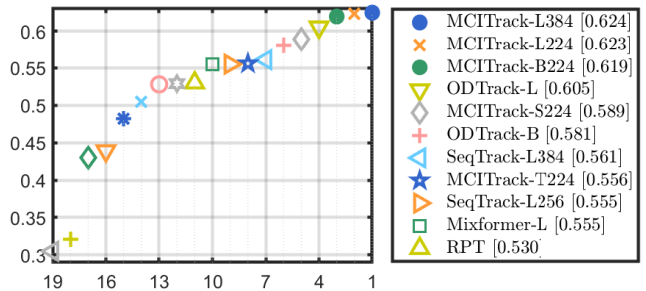


Figure 4: EAO rank plots on VOT2020.

shown in Table 1, MCITrack-B224 achieves the highest AO score of 77.9%, surpassing the second-place ODTrack-B384 by 0.9%. Additionally, among the large-scale models, MCITrack-L384 achieves the top AO score of 80.0%.

TNL2K, NFS and UAV123. We also evaluate our MCITrack models on three additional datasets: TNL2K (Wang et al. 2021b), NFS (Kiani Galoogahi et al. 2017), and UAV123 (Mueller, Smith, and Ghanem 2016). TNL2K is a recently released large-scale dataset with 700 video sequences, while NFS and UAV123 are two smaller benchmarks, containing 100 and 123 videos, respectively. As shown in Table 3, MCITrack achieves new state-of-the-art performance on TNL2K and NFS, with AUC scores of 65.3% and 71.1%, respectively. On UAV123, MCITrack-L384 also delivers a competitive result, achieving an AUC score of 71.5%, which is better than most previous trackers.

VOT2020. VOT2020 (Kristan et al. 2020) contains 60 challenging videos. We equip MCITrack with AlphaRefine (Yan et al. 2021b) to predict segmentation masks. As shown in Figure 4, MCITrack-L384, MCITrack-L224, and MCITrack-B224 secure the top three positions with EAO scores of 62.4%, 62.3%, and 61.9%, respectively.

Ablation Study

In the ablation study, we use MCITrack-B224 as our baseline model to investigate the impact of the CIF module, video clip length, and different methods of contextual information transmission on the model.

CIF Module. The CIF module is a crucial component of MCITrack, responsible for storing and transmitting contextual information. We conduct ablation studies on the number of CIF blocks in the CIF module, the structure of the CIF block, and the size of the hidden state. The results are shown in Table 4. Our baseline model uses 4 CIF blocks with a hidden state size of 16. As observed (#2 and #3), decreasing or increasing the number of CIF blocks to 2 or 6 results in a reduction in AUC on LaSOT by 1.4% and 0.5%, respectively. This demonstrates the effectiveness of our method for deep integration with the backbone, while too many interactions increase model complexity and can degrade performance. We also explore different structures of the CIF block, #4, #5, and #6 show the results. By replacing the attention modules with addition, we find that removing in-attention, out-

#	Method	AUC	P_{Norm}	P	Δ
1	Baseline	75.3	85.6	83.3	-
2	2 CIF blocks	73.9	84.5	81.8	-1.4
3	6 CIF blocks	74.8	84.8	82.8	-0.5
4	Without in cross attention	74.3	84.3	81.8	-1.0
5	Without out cross attention	74.2	84.1	82.0	-1.1
6	Without in & out cross attention	73.5	83.3	80.7	-1.8
7	4 hidden states	74.0	83.9	81.7	-1.3
8	8 hidden states	74.6	84.8	82.7	-0.7
9	32 hidden states	74.3	84.2	81.9	-1.0

Table 4: Ablation Study of the CIF module on LaSOT. #1 denotes the baseline setting. #2 and #3 denote the number of CIF blocks. #4, #5, and #6 denote the CIF block structure. #7, #8, and #9 denote the hidden state size. Δ denotes the performance change (AUC) compared with the baseline.

Video Clip Length	AUC	P_{Norm}	P	Δ
Baseline	75.3	85.6	83.3	-
2 frames	72.7	82.4	80.0	-2.6
3 frames	73.8	83.5	81.3	-1.5
4 frames	74.4	84.2	81.9	-0.9
6 frames	73.8	83.8	81.3	-1.5

Table 5: Ablation Study of the video clip length on LaSOT. Δ denotes the performance change (AUC) compared with the baseline.

attention, or both reduces the model performance by 1.0%, 1.1%, and 1.8%, respectively. These results underscore the effectiveness of our CIF block structure. Table 4 (#7, #8, and #9) also demonstrates the impact of different hidden state sizes. Both too small and too large hidden state sizes lead to decreased model performance, indicating the importance of optimal hidden state size for maintaining performance.

Video Clip Length. Table 5 shows the impact of video clip length on MCITrack. Our baseline uses a 5-frame clip. The results reveal that as the video clip length increases, incorporating more appearance information, the model performance also improves, with an increase of 2.6% in AUC on LaSOT when extending from 2 frames to 5 frames. However, further increasing the video clip length does not consistently lead to performance gains and may even cause performance degradation. For instance, using 6 frames results in a 1.5% decrease in AUC on LaSOT. This suggests that excessively long video clips can impose a higher learning burden on the model, potentially leading to reduced performance.

Contextual Information Propagation Methods. Table 6 compares different methods for transmitting contextual information. The baseline is MCITrack, which uses Mamba’s hidden states to transmit contextual information. ”Wo CI” indicates the absence of contextual information in tracking. ”Extra Token” refers to use additional tokens for transmitting contextual information. ”LSTM” indicates use LSTM’s hidden states for this purpose, while ”Previous Features” involves using features from the previous frame directly. Compared to the baseline, ”Wo CI” results in a 2.3% decrease

Method	AUC	P_{Norm}	P	Δ
Baseline	75.3	85.6	83.3	-
Wo CI	73.0	83.0	80.6	-2.3
Extra Token	73.4	83.0	80.7	-1.9
Previous Features	73.4	83.2	80.6	-1.9
LSTM	74.3	84.2	81.8	-1.0

Table 6: Ablation Study of different contextual information propagation methods on LaSOT. Δ denotes the performance change (AUC) compared with the baseline.

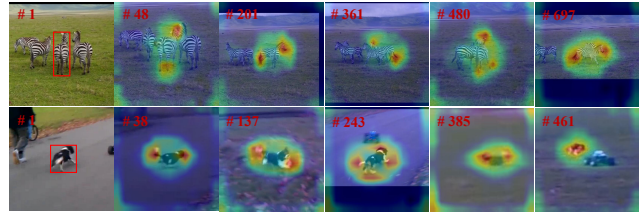


Figure 5: Visualization of the output features of CIF module.

in AUC on LaSOT, highlighting the importance of contextual information for the tracker. ”Extra Token” and ”Previous Features” result in a 1.9% decrease in AUC compared to the baseline. In contrast, ”LSTM” achieves a relatively higher AUC of 74.3%, which is 0.9% better than ”Extra Token” and ”Previous Features”. This suggests that our method of using hidden states to transmit contextual information is effective. Compared with previous methods, our method can convey richer and more crucial contextual information.

Visualization. Figure 5 presents the visualization results of the output features from the CIF module. It shows that, after passing through the CIF module, the model emphasizes critical edge information of the target, such as the tail, legs, and head of the zebra and the dog. This enhancement allows for more accurate target localization.

Conclusion

In this work, we introduce a new family of video-level tracking models named MCITrack, which develops a new method for transmitting contextual information. Its core component is the Contextual Information Fusion(CIF) module, which utilizes hidden states from Mamba to effectively transmit richer and more crucial contextual information. Experiments show that MCITrack is effective, achieving state-of-the-art performance across multiple datasets. We hope this work inspires further research in video-level tracking.

Limitation. Despite its significant results, MCITrack faces challenges with slow model training due to its video-level modeling approach. Additionally, handling the video clip introduces more computational overhead compared to the single-frame image. To address these issues, it is essential to accelerate model training and minimize the computational burden of the video clip. A promising solution is to efficiently integrate multi-frame image information to reduce the video clip size and develop a lightweight model.

Acknowledgements

The paper is supported in part by National Natural Science Foundation of China (nos. U23A20384 and 62476044), in part by Key Research and Development Program of Liaoning Province (no. 2023JH26/10200015), and in part by Open Research Fund of the National Key Laboratory of Space Intelligent Control Technology (No. HTKJ2024KL502016).

References

- Bai, Y.; Zhao, Z.; Gong, Y.; and Wei, X. 2024. ARTrackV2: Prompting Autoregressive Tracker Where to Look and How to Describe. In *CVPR*, 19048–19057.
- Bertinetto, L.; Valmadre, J.; Henriques, J. F.; Vedaldi, A.; and Torr, P. H. S. 2016. Fully-Convolutional Siamese Networks for Object Tracking. In *ECCV*, 850–865.
- Bhat, G.; Danelljan, M.; Gool, L. V.; and Timofte, R. 2019. Learning Discriminative Model Prediction for Tracking. In *ICCV*, 6182–6191.
- Cai, Y.; Liu, J.; Tang, J.; and Wu, G. 2023. Robust Object Modeling for Visual Tracking. In *ICCV*, 9589–9600.
- Cao, Z.; Huang, Z.; Pan, L.; Zhang, S.; Liu, Z.; and Fu, C. 2022. TCTrack: Temporal Contexts for Aerial Tracking. In *CVPR*, 14798–14808.
- Chang, L.; Yongsheng, Y.; Xin, C.; Huchuan, L.; and Dong, W. 2024. Spatial-temporal initialization dilemma: towards realistic visual tracking. *Visual Intelligence 2, Article no. 35*.
- Chen, B.; Li, P.; Bai, L.; Qiao, L.; Shen, Q.; Li, B.; Gan, W.; Wu, W.; and Ouyang, W. 2022. Backbone is All Your Need: A Simplified Architecture for Visual Object Tracking. In *ECCV*, 375–392.
- Chen, X.; Peng, H.; Wang, D.; Lu, H.; and Hu, H. 2023a. SeqTrack: Sequence to Sequence Learning for Visual Object Tracking. In *CVPR*, 14572–14581.
- Chen, X.; Yan, B.; Zhu, J.; Lu, H.; Ruan, X.; and Wang, D. 2023b. High-Performance Transformer Tracking. *IEEE TPAMI*, 8507–8523.
- Chen, X.; Yan, B.; Zhu, J.; Wang, D.; Yang, X.; and Lu, H. 2021. Transformer Tracking. In *CVPR*, 8126–8135.
- Chen, Z.; Zhong, B.; Li, G.; Zhang, S.; and Ji, R. 2020. Siamese Box Adaptive Network for Visual Tracking. In *CVPR*, 6668–6677.
- Cui, Y.; Jiang, C.; Wang, L.; and Wu, G. 2022. MixFormer: End-to-End Tracking with Iterative Mixed Attention. In *CVPR*, 13608–13618.
- Cui, Y.; Jiang, C.; Wang, L.; and Wu, G. 2024. MixFormer: End-to-End Tracking with Iterative Mixed Attention. *IEEE TPAMI*, 0–18.
- Danelljan, M.; Bhat, G.; Khan, F. S.; and Felsberg, M. 2019. ATOM: Accurate Tracking by Overlap Maximization. In *CVPR*, 4660–4669.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*, 1–9.
- Elfwing, S.; Uchibe, E.; and Doya, K. 2018. Sigmoid-Weighted Linear Units for Neural Network Function Approximation in Reinforcement Learning. *Neural Networks*, 3–11.
- Fan, H.; Bai, H.; Lin, L.; Yang, F.; Chu, P.; Deng, G.; Yu, S.; Huang, M.; Liu, J.; Xu, Y.; et al. 2021. LaSOT: A High-Quality Large-Scale Single Object Tracking Benchmark. *IJCV*, 439–461.
- Fan, H.; Lin, L.; Yang, F.; Chu, P.; Deng, G.; Yu, S.; Bai, H.; Xu, Y.; Liao, C.; and Ling, H. 2019. LaSOT: A High-Quality Benchmark for Large-Scale Single Object Tracking. In *CVPR*, 5374–5383.
- Fu, Z.; Liu, Q.; Fu, Z.; and Wang, Y. 2021. STMTrack: Template-free Visual Tracking with Space-time Memory Networks. In *CVPR*, 13774–13783.
- Gao, S.; Zhou, C.; Ma, C.; Wang, X.; and Yuan, J. 2022. AiATrack: Attention in Attention for Transformer Visual Tracking. In *ECCV*, 146–164.
- Graves, A. 2012. *Supervised Sequence Labelling with Recurrent Neural Networks*. Studies in Computational Intelligence. Springer.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-Time Sequence Modeling with Selective State Spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, A.; Goel, K.; and Ré, C. 2022. Efficiently Modeling Long Sequences with Structured State Spaces. In *ICLR*, 1–9.
- Gu, A.; Johnson, I.; Goel, K.; Saab, K.; Dao, T.; Rudra, A.; and Ré, C. 2021. Combining Recurrent, Convolutional, and Continuous-Time Models with Linear State Space Layers. In *NeurIPS*, 572–585.
- Guo, M.; Zhang, Z.; Fan, H.; Jing, L.; Lyu, Y.; Li, B.; and Hu, W. 2022. Learning Target-aware Representation for Visual Tracking via Informative Interactions. In *IJCAI*, 927–934.
- He, K.; Zhang, C.; Xie, S.; Li, Z.; and Wang, Z. 2023. Target-Aware Tracking with Long-Term Context Attention. In *AAAI*, 773–780.
- Huang, L.; Zhao, X.; and Huang, K. 2019. GOT-10k: A Large High-Diversity Benchmark for Generic Object Tracking in The Wild. *IEEE TPAMI*, 1562–1577.
- Kiani Galoogahi, H.; Fagg, A.; Huang, C.; Ramanan, D.; and Lucey, S. 2017. Need for Speed: A Benchmark for Higher Frame Rate Object Tracking. In *ICCV*, 1125–1134.
- Kristan, M.; Leonardis, A.; Matas, J.; Felsberg, M.; Pflugfelder, R.; Kämäräinen, J.-K.; Danelljan, M.; Zajc, L. Č.; Lukežič, A.; Drbohlav, O.; et al. 2020. The Eighth Visual Object Tracking VOT2020 Challenge Results. In *ECCV*, 547–601.
- Li, B.; Wu, W.; Wang, Q.; Zhang, F.; Xing, J.; and Yan, J. 2019. SiamRPN++: Evolution of Siamese Visual Tracking with Very Deep Networks. In *CVPR*, 4282–4291.
- Li, B.; Yan, J.; Wu, W.; Zhu, Z.; and Hu, X. 2018. High Performance Visual Tracking With Siamese Region Proposal Network. In *CVPR*, 8971–8980.

- Li, X.; Huang, Y.; He, Z.; Wang, Y.; Lu, H.; and Yang, M.-H. 2023. Citetracker: Correlating Image and Text for Visual Tracking. In *ICCV*, 9974–9983.
- Lin, L.; Fan, H.; Xu, Y.; and Ling, H. 2022. Swintrack: A Simple and Strong Baseline for Transformer Tracking. In *NeurIPS*, 16743–16754.
- Lin, L.; Fan, H.; Zhang, Z.; Wang, Y.; Xu, Y.; and Ling, H. 2024. Tracking Meets LoRA: Faster Training, Larger Model, Stronger Performance. In *ECCV*, 1–15.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *ICCV*, 2980–2988.
- Lin, T.-Y.; Maire, M.; Belongie, S. J.; Bourdev, L. D.; Girshick, R. B.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In *ECCV*, 740–755.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; and Liu, Y. 2024. VMamba: Visual State Space Model. *arXiv preprint arXiv:2401.10166*.
- Loshchilov, I.; and Hutter, F. 2018. Decoupled Weight Decay Regularization. In *ICLR*, 1–9.
- Mayer, C.; Danelljan, M.; Bhat, G.; Paul, M.; Paudel, D. P.; Yu, F.; and Van Gool, L. 2022. Transforming Model Prediction for Tracking. In *CVPR*, 8731–8740.
- Mayer, C.; Danelljan, M.; Paudel, D. P.; and Van Gool, L. 2021. Learning Target Candidate Association to Keep Track of What Not to Track. In *ICCV*, 13444–13454.
- Mueller, M.; Smith, N.; and Ghanem, B. 2016. A Benchmark and Simulator for UAV Tracking. In *ECCV*, 445–461.
- Muller, M.; Bibi, A.; Giancola, S.; Alsubaihi, S.; and Ghanem, B. 2018. TrackingNet: A Large-Scale Dataset and Benchmark for Object Tracking in The Wild. In *ECCV*, 300–317.
- Peng, L.; Gao, J.; Liu, X.; Li, W.; Dong, S.; Zhang, Z.; Fan, H.; and Zhang, L. 2024. VastTrack: Vast Category Visual Object Tracking. *arXiv preprint arXiv:2403.03493*.
- Rezatofghi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I. D.; and Savarese, S. 2019. Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In *CVPR*, 658–666.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI*, 234–241.
- Shi, L.; Zhong, B.; Liang, Q.; Li, N.; Zhang, S.; and Li, X. 2024. Explicit Visual Prompts for Visual Object Tracking. In *AAAI*, 4838–4846.
- Song, Z.; Luo, R.; Yu, J.; Chen, Y.-P. P.; and Yang, W. 2023. Compact Transformer Tracker with Correlative Masked Modeling. In *AAAI*, 2321–2329.
- Tao, R.; Gavves, E.; and Smeulders, A. W. M. 2016. Siamese Instance Search for Tracking. In *CVPR*, 1420–1429.
- Tian, Y.; Xie, L.; Qiu, J.; Jiao, J.; Wang, Y.; Tian, Q.; and Ye, Q. 2024. Fast-iTPN: Integrally Pre-Trained Transformer Pyramid Network with Token Migration. *IEEE TPAMI*, 1–15.
- Voigtlaender, P.; Luiten, J.; Torr, P. H. S.; and Leibe, B. 2020. Siam R-CNN: Visual Tracking by Re-Detection. In *CVPR*, 6578–6588.
- Wang, N.; Zhou, W.; Wang, J.; and Li, H. 2021a. Transformer Meets Tracker: Exploiting Temporal Context for Robust Visual Tracking. In *CVPR*, 1571–1580.
- Wang, X.; Shu, X.; Zhang, Z.; Jiang, B.; Wang, Y.; Tian, Y.; and Wu, F. 2021b. Towards More Flexible and Accurate Object Tracking with Natural Language: Algorithms and Benchmark. In *CVPR*, 13763–13773.
- Wang, Z.; Zheng, J.-Q.; Zhang, Y.; Cui, G.; and Li, L. 2024. Mamba-Unet: Unet-Like Pure Visual Mamba for Medical Image Segmentation. *arXiv preprint arXiv:2402.05079*.
- Wei, X.; Bai, Y.; Zheng, Y.; Shi, D.; and Gong, Y. 2023. Autoregressive Visual Tracking. In *CVPR*, 9697–9706.
- Xie, F.; Chu, L.; Li, J.; Lu, Y.; and Ma, C. 2023. VideoTrack: Learning to Track Objects via Video Transformer. In *CVPR*, 22826–22835.
- Xie, F.; Wang, C.; Wang, G.; Cao, Y.; Yang, W.; and Zeng, W. 2022. Correlation-Aware Deep Tracking. In *CVPR*, 8751–8760.
- Xie, J.; Zhong, B.; Mo, Z.; Zhang, S.; Shi, L.; Song, S.; and Ji, R. 2024. Autoregressive Queries for Adaptive Tracking with Spatio-Temporal Transformers. In *CVPR*, 19300–19309.
- Xu, Y.; Wang, Z.; Li, Z.; Yuan, Y.; and Yu, G. 2020. SiamFC++: Towards Robust and Accurate Visual Tracking with Target Estimation Guidelines. In *AAAI*, 12549–12556.
- Yan, B.; Peng, H.; Fu, J.; Wang, D.; and Lu, H. 2021a. Learning Spatio-Temporal Transformer for Visual Tracking. In *ICCV*, 10448–10457.
- Yan, B.; Zhang, X.; Wang, D.; Lu, H.; and Yang, X. 2021b. Alpha-Refine: Boosting Tracking Performance by Precise Bounding Box Estimation. In *CVPR*, 5289–5298.
- Ye, B.; Chang, H.; Ma, B.; Shan, S.; and Chen, X. 2022. Joint Feature Learning and Relation Modeling for Tracking: A One-Stream Framework. In *ECCV*, 341–357.
- Zhang, B.; and Sennrich, R. 2019. Root Mean Square Layer Normalization. In *NeurIPS*, 12360–12371.
- Zhang, Z.; and Peng, H. 2019. Deeper and Wider Siamese Networks for Real-Time Visual Tracking. In *CVPR*, 4591–4600.
- Zhang, Z.; Peng, H.; Fu, J.; Li, B.; and Hu, W. 2020. Ocean: Object-aware Anchor-free Tracking. In *ECCV*, 771–787.
- Zheng, Y.; Zhong, B.; Liang, Q.; Mo, Z.; Zhang, S.; and Li, X. 2024. ODTrack: Online Dense Temporal Token Learning for Visual Tracking. In *AAAI*, 7588–7596.
- Zhu, J.; Lai, S.; Chen, X.; Wang, D.; and Lu, H. 2023. Visual prompt multi-modal tracking. In *CVPR*, 9516–9526.
- Zhu, L.; Liao, B.; Zhang, Q.; Wang, X.; Liu, W.; and Wang, X. 2024. Vision Mamba: Efficient Visual Representation Learning with Bidirectional State Space Model. *arXiv preprint arXiv:2401.09417*.