

Empowering Embodied Visual Tracking with Visual Foundation Models and Offline RL

Fangwei Zhong^{1,2*}, Kui Wu^{3*}, Hai Ci⁴, Churan Wang¹, and Hao Chen⁵

¹ Peking University, Beijing, China

² State Key Laboratory of General Artificial Intelligence, BIGAI&PKU, Beijing, China

³ State Key Laboratory of Complex & Critical Software Environment, Beihang University, Beijing, China

⁴ National University of Singapore, Singapore

⁵ City University of Macau, Macao, China

{zfw, cihai, churanwang}@pku.edu.cn {wukui0099, sundaychenhao}@gmail.com

Abstract. Embodied visual tracking is to follow a target object in dynamic 3D environments using an agent’s egocentric vision. This is a vital and challenging skill for embodied agents. However, existing methods suffer from inefficient training and poor generalization. In this paper, we propose a novel framework that combines visual foundation models (VFM) and offline reinforcement learning (offline RL) to empower embodied visual tracking. We use a pre-trained VFM, such as “Tracking Anything”, to extract semantic segmentation masks with text prompts. We then train a recurrent policy network with offline RL, e.g., Conservative Q-Learning, to learn from the collected demonstrations without online interactions. To further improve the robustness and generalization of the policy network, we also introduce a mask re-targeting mechanism and a multi-level data collection strategy. In this way, we can train a robust policy within an hour on a consumer-level GPU, e.g., Nvidia RTX 3090. We evaluate our agent on several high-fidelity environments with challenging situations, such as distraction and occlusion. The results show that our agent outperforms state-of-the-art methods in terms of sample efficiency, robustness to distractors, and generalization to unseen scenarios and targets. We also demonstrate the transferability of the learned agent from virtual environments to a real-world robot. ⁶

1 Introduction

Embodied agents, such as robots or virtual avatars, need to track their interested objects (target) in their surroundings with visual observation while performing social tasks. For example, a robot may need to follow a user’s path to prepare for help. Embodied visual tracking (EVT) has a wide range of applications, including drones [8], mobile robots [32], and unmanned vehicles [15, 16].

* indicates equal contribution

⁶ Project Website: <https://sites.google.com/view/offline-evt>




Fig. 1: The overall framework of our proposed method. For data acquisition, a state-based policy is employed to collect diverse image-action-reward trajectories through interactions with various complex environments. We augment the policy by adding different levels of perturbation in actions. The observed color images are encoded into text-conditioned segmentation masks, which highlight the target object (white), show the obstacle (colorful), and remove the background noise (black). Subsequently, we employ offline reinforcement learning, such as conservative Q-learning (CQL), to train the recurrent policy network, which outputs actions based on the segmentation masks.

However, EVT faces several challenges in applications: 1) **Training Efficiency:** Training the tracking policy by Reinforcement Learning (RL) is a common approach [24, 42], but it is computationally intensive and time-consuming. The agent requires extensive trial-and-error interactions with the environment for learning, often taking more than 12 hours. 2) **Cross-Domain Generalization:** Agents are required to track unseen targets with different visual appearances and motion patterns in diverse environments. Thus, it faces a large domain gap between the training and testing environments. 3) **Spatial and Temporal Reasoning:** When tracking in complex environments, the agent must infer the dynamics of the scene spatially and temporally to handle obstacles, occlusions, and distractors. 4) **Real-Time Interaction:** As the target moves, the embodied agent must react quickly with limited computational resources. This imposes challenges in balancing efficiency and accuracy at inference time.

In this paper, we propose an efficient training framework for embodied visual tracking, by integrating visual foundation models and offline reinforcement learning. This framework stands out as an efficient method for training a visual policy. Our framework consists of four main components: 1) A data acquisition pipeline can generate multi-level imperfect state-based expert policy to automatically collect diverse trajectories in an augmentable virtual environment. 2) A text-conditioned segmentation mask, provided as an abstracted representation of the target and obstacles. 3) A recurrent policy network to filter the noise

in the masks and output temporal consistent actions for tracking. 4) An offline reinforcement learning algorithm trains the policy network with the collected data. Within this framework, we train a robust tracking policy within an hour on a consumer-level GPU, e.g., Nvidia RTX 3090. We evaluate our tracker on a set of high-fidelity virtual environments that simulate diverse and challenging situations, such as active distraction, occlusion, and reflection. The results show that our framework achieves superior performance in terms of accuracy, robustness, and generalization. We also validate the sim2real transferability by deploying our method in a real-world robot.

Our main contributions are in threefolds: 1) We propose a framework that incorporates the visual foundation models and offline reinforcement learning with practical techniques, such as text-conditioned segmentation mask, re-targeting mechanism, and recurrent policy, to efficiently build a robust embodied visual tracking agent. 2) We introduce a data acquisition method with multi-level policy perturbation and collect an offline dataset for offline RL. 3) We evaluate the proposed method in both virtual and real environments, showing its significant improvements in robustness and generalization.

2 Related work

Embodied Visual Tracking (EVT) is an active research area in embodied vision [9, 22, 24, 25, 33, 39–43]. Luo, et al [24] first trained an end-to-end visual policy via RL and successfully deployed the policy trained in a simulator on a real-world robot [25]. Zhong, et al. [41, 42] use multi-agent games to improve the robustness and generalization of the tracker in unseen environments. Recently, Meta has adopted a similar end-to-end RL method for embodied social navigation tasks [27]. Zhong, et al. [43] and Bajcsy, et al. [2, 14] use cross-modal teacher-student learning strategy with multi-agent games to train a robust visual tracking policy. However, all these methods are trained by online interactions, leading to a high cost in training time. In addition, training time might increase when transferring to complex real-world environments, such as obstacle occlusions, active distractors, and temporary targets. To improve training efficiency, we first introduce an offline reinforcement learning paradigm for embodied visual tracking, making it feasible to efficiently train a robust tracker in an hour.

Offline Reinforcement Learning (Offline RL) [10, 21] has been first proposed as a paradigm that learns optimal policies from fixed offline datasets, avoiding the need for online data collection, which aims to increase the training efficiency. Shah, et al. [31] presented the first offline RL system for robotic navigation, which uses previously collected data to optimize user-specified reward functions in the real world. Yet, this system depends on a fixed reward function, making it less adaptable to various environments. With superior training efficiency, Offline RL still faces several challenges, such as data distributional shifts, inaccurate state reward estimation, and exploration-exploitation trade-offs.

Visual Reinforcement Learning (Visual RL) faces challenges due to the high-dimensional and noisy observation space. End-to-end methods typically en-

code input images directly into a latent space and often suffer from inefficiency and poor generalization. To address these issues, several works have explored the impact of mid-level representations on final performance [26, 30, 35, 37, 38, 44], indicating that while mid-level representations generalize well across domains, the optimal representation varies by task. For instance, CSR [11] extracting bounding boxes and masks to create an object state graph, which enhances spatial awareness of visual representation for embodied agent navigation tasks [5]. RSPT [40] introduces a motion-aware structure representation for tracking in cluttered environments. In this paper, we leverage visual foundation models to generate text-conditioned object masks as state representations. The learned policy can generalize well to unseen environments and targets. Our solution may also inspire other visual reinforcement learning tasks.

Visual Foundation Models (VFMs) [1, 4] learn versatile visual representations from large-scale data for tasks like scene understanding, object detection, and segmentation. However, in embodied vision, where agents interact with dynamic environments, VFMs face challenges such as lack of spatial awareness and robustness. For instance, CLIP [29] struggles with precise localization, and SAM [19] and Grounding-DINO [23] are not robust to noise, occlusion, and motion. To overcome these limitations, models like TAM [36], E-SAM [34], DEVA [6], and SAM-Track [7] integrate VFMs with other models to improve segmentation and tracking. In robotic tasks, VFMs are used for state representations or goal states, as seen in EmbCLIP [18] for navigation and DALL-E-Bot [17] for object rearrangement. LfVoid [12] links instructions and visual goals in reinforcement learning. In this paper, we leverage text-based video object segmentation models to generate text-conditioned segmentation masks for embodied visual tracking, enhancing spatial and temporal understanding and improving policy generalization with a re-targeting mechanism.

3 Task Definition

Embodied visual tracking is a typical embodied vision task that involves both visual perception and movement control. The tracking agent needs to cope with occlusions, distractors, and high-dynamic environments. The task can be formulated as a Partial Observable Markov Decision Process (POMDP), with state $s \in \mathcal{S}$, partial observation $o \in \mathcal{O}$, action $a \in \mathcal{A}$, and reward $r \in \mathcal{R}$. At each time step t , the agent observes an image o_t captured by its embodied vision system and outputs an action a_t to maintain a relative position to the target object. The agent receives a reward $r(s_t)$ that reflects how well it tracks the target. In the agent-centric coordinate system, the reward function is defined as $r = 1 - \frac{|\rho - \rho^*|}{\rho_{max}} - \frac{|\theta - \theta^*|}{\theta_{max}}$, where (ρ, θ) denotes the current target position relative to the tracker, $(\rho^*, \theta^*) = (2.5m, 0)$ represents the expected target position, i.e., the target should be $2.5m$ in front of the tracker. The error is normalized by the field of the view $(\rho_{max}, \theta_{max})$. This reward function is similar to the one used in [40, 41]. The goal of the agent is to learn a policy $\pi(a_t|s_t)$ that maximizes the expected cumulative reward over a finite horizon T , i.e., $E_\pi \left[\sum_{t=0}^T r_t \right]$.

In this paper, we consider the offline training setting that aims to learn a policy with a fixed dataset of transitions $\mathcal{D} = \{(\mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_T)_i | i = 1 \dots N\}$, where T is the episode length, N is the number of episodes. A transition is a tuple of state, action, reward, and next state, denoted as $\mathcal{T}_t = (s_t, a_t, r_t, s_{t+1})$. In this setting, the agent is not required any online interaction with the environment during training. This is crucial because online interaction has been identified as a primary bottleneck in enhancing the efficiency of the training process. In this setting, the main challenge is to find an effective solution to improve the generalization of the policy to handle the unseen state transition that is not contained in the offline dataset.

4 Method

In this section, we introduce our framework to learn visual embodied tracking in an offline manner. First, we employ the visual foundation models, e.g., DEVA [6] or SAM-Track [7], to build a generalizable text-conditioned segmentation mask as state representation and a multi-level expert generation strategy to collect diverse demonstrations. At the end, we train a recurrent policy network via offline reinforcement learning, e.g., Conservative Q-learning [21].

4.1 Text-Conditioned Segmentation Mask

We find that vanilla end-to-end offline reinforcement learning does not work well for embodied visual tracking (EVT) due to the high-dimensional visual observation and the complex dynamic environment. We propose a text-conditioned segmentation mask as the state representation for EVT, which should be spatial-temporal consistent, domain-invariant, and efficient for real-time execution. We use an open-vocabulary segmentation model and a re-targeting mechanism to obtain a temporal-consistent text-conditioned segmentation mask. The open-vocabulary tracking model can abstract and generalize the state representation by combining semantic knowledge and visual features. This representation can be provided by SAM-Track [7], DEVA [6], and other open-vocabulary segmentation models utilizing text prompts. The re-targeting mechanism can make the color of the target in the mask consistent across different episodes and enhance the contrast between the target and the environment, improving the robustness to visual distraction. We introduce the details of the two steps as follows:

1) **Segmentation with Text Prompts:** We use the text prompt to specify the category of the desired target object and the related background objects for segmentation. For example, when tracking a person in a cluttered environment, we can specify the prompt as $\{“person”, “obstacles”\}$, thereby drawing a segmentation mask identifying each person and obstacle uniquely. In contrast to Instance Segmentation Mask (ISM), Text-Conditioned Segmentation Mask can effectively eliminate extraneous background details, thus reducing the presence of irrelevant noise and mitigating the domain gap. If we only segment




Fig. 2: Examples of an observed RGB image and the corresponding instance segmentation mask (ISM) from UnrealCV [28] and Text-conditioned segmentation mask from DEVA. The middle one is the original mask from DEVA [6]. The two right masks utilize re-targeting to emphasize the target with different text prompts.

the target, the obstacle will be missed in the representation, leading to collision in complex environments. Specifically, at time step t , the VFM with text prompt ϕ_{prompt} extracts a segmentation mask S_t from the image O_t , expressed as $S_t = \phi_{prompt}(O_t)$. The mask contains K instances, each denoted as $m(i, c)$, where $m(i, c) \in S_t$, c is the object category $c \in \{“person”, “obstacles”\}$ and i is the instance ID. Figure 2 shows the ISM and text-conditioned masks in different configurations.

2) **Re-targeting Mechanism:** We can not directly identify the target from the original masks provided by the VFMs, as they randomly assign different colors to each instance. Such masks will require our policy to learn to reason the target from a sequence of masks. To explicitly identify the target in the mask, we introduce a re-targeting mechanism. At the initial frame of each episode, we specify the initial target mask using a scoring system that prioritizes saliency, based on size and centrality. Masks covering over 15% of the image are scored with: $Score = s_{m(i,c)} * (1 - |e_{pos}|)$, where $s_{m(i,c)}$ is the mask’s area size and e_{pos} is the normalized position error from the mask center to the image center. The mask with the highest score is selected as the target object $mask_{(i,c)}$, ensuring that it is large and centrally located. We record the index i as the target ID. In the subsequent propagation procedure, all object masks associated with the target ID have their RGB values manually adjusted to (255, 255, 255), resulting in a white color representation, as shown in Figure 1.

4.2 Multi-level Demonstration Collection

For offline policy learning, the diversity of the demonstration determines the skill level of the learned policy. However, in many cases, it is unrealistic and inappropriate to collect training data by deploying a real robot because of the high cost of trial-and-error interactions. Besides, the accurate reward signal is necessary for learning, which is hard to obtain in real-world scenarios. We use an augmented virtual environment to automatically collect large-scale tracking demonstrations for offline policy learning. By leveraging UnrealCV [28], we control environmental factors to enhance diversity and access ground truth states for building an expert for EVT. Our data collection process involves three steps:

1) **Environment Randomization:** We stochastically initialize various environmental factors, including lighting conditions, the number of obstacles and players, as well as the positions, rotations, and textures of these elements. This configuration intentionally includes challenging scenarios, such as the presence of multiple distracting agents and unknown environmental structures. We randomly select a tracker and a target from the pool of all players. Subsequently, the expected location of the tracker is positioned $2.50m$ behind the target. The tracker’s expected rotation angle is carefully adjusted to maintain the target at the center of the field of view, as shown in Figure 2.

2) **Multi-Level Data Generation:** We develop a navigation system in Unreal Engine to generate diverse trajectories for the target and distractors. The system enables agents to autonomously navigate to any reachable locations via the shortest path. For tracking, we use a state-based PID controller as the *expert policy*. To simulate the policy with different skill levels, we add noise to the output actions. We define a probability threshold μ for switching to random actions for a random number of continuous steps. By adjusting μ , we can generate multi-level demonstrations, from random walking to closed following. We apply the state-based tracking policy for 500 steps per episode, and record the tracker’s visual observations, rewards, actions, and done flags of each time step. Each data element is a tuple of $(o_t, a_t, r_t, done)$, where o_t is the observation (first-person view RGB image of the tracker), as shown in Figure 2, a_t is the action taken by the agent, r_t is the reward of the current step, and *done* is a boolean indicating the end of the episode. The data elements are ordered by their time stamps.

3) **Pre-Processing Raw Observation:** Before training, we use the visual foundation model and re-targeting mechanism to transform all raw-pixel images to text-conditioned segmentation masks, as shown in Figure 2. Notably, the re-targeting mechanism is a simple but effective method to highlight the target from other objects in the segmentation mask (the target is shown in white).

4.3 Learning Recurrent Policy via Offline RL

The recurrent policy network contains a segmentation mask encoder, a temporal encoder, and an actor-critic network. We train it using offline reinforcement learning with the collected dataset, as illustrated in Figure 1.

Recurrent Policy Networks: Drawing from prior experience, the agent may encounter challenges related to partial observability and non-Markovian movements [43]. For instance, abrupt changes in speed or direction near obstacles could lead to the loss of the target within the field of view. Traditional single-frame state representations lack temporal information and are ill-equipped to handle such scenarios during long-horizon interactions with the environment. To address these issues, we adopt a standard CNN-LSTM architecture as our recurrent state encoder, which effectively captures long-term hidden state representations and also maintains a real-time inference. In offline training, we first sample a batch sequence of data from the collected offline dataset as input, each sequence of data is expressed as $\{(S_i, a_i, S'_i, r_i) \dots (S_{i+L}, a_{i+L}, S'_{i+L}, r_{i+L})\}$ where i is a randomly selected starting frame in an offline dataset. S'_i is the next state

of S_i . L is the predefined sequence length. In the training phase, the sequence of the pre-processed segmentation masks $S_i, S_{i+1} \dots S_{i+L}$ are encoded into state representation $x_i, x_{i+1} \dots x_t$, which is used in Q network to learn offline action policy π_{off} . In the testing phase, the current segmentation result S_t and the last hidden state h_{t-1}, c_{t-1} are used to get the action a_t .

Conservative Q-learning with SAC: To optimize the policy network, we employ the Conservative Q-learning (CQL) [21] with the Soft Actor-Critic (SAC) networks [13] for our methods, referred to as CQL-SAC. CQL-SAC consists of three neural networks: two critic network Q_θ^1, Q_θ^2 , used for estimate Q values and an actor network V_ϕ , used for learning control policy. θ and ϕ are network parameters. Moreover, CQL-SAC uses an entropy regularization coefficient α , which controls the degree of exploration. The Q-functions are updated by minimizing the following objective function:

$$L_\theta(D) = \mathbb{E}_{\mathbf{s} \sim \mathcal{D}} \left[\log \sum_{\mathbf{a}} \exp(Q_\theta^i(\mathbf{s}, \mathbf{a})) - \mathbb{E}_{\mathbf{a} \sim \pi_\phi(\mathbf{a}|\mathbf{s})} [Q_\theta^i(\mathbf{s}, \mathbf{a})] \right] + \frac{1}{2} \mathbb{E}_{\mathbf{s}, \mathbf{a}, \mathbf{s}' \sim \mathcal{D}} \left[\left(Q_\theta^i(\mathbf{s}, \mathbf{a}) - (r + \gamma E_{\mathbf{a}' \sim \pi_\phi} [Q_{min}(s', a') - \alpha \log \pi_\phi(a' | s')]) \right)^2 \right] \quad (1)$$

where $i \in \{1, 2\}$, $Q_{min} = \min_{i \in \{1, 2\}} Q_\theta^i$, γ is the discount factor, π_ϕ is the policy that derived from the actor network V_ϕ . The first term is the CQL regularization term, which encourages the Q-function to be close to the log-density of the data distribution, and the second term is the standard SAC objective.

5 Experiments

In this section, we conduct comprehensive experiments to validate: 1) the training efficiency, 2) the generalization ability to unseen environments and unseen target categories, 3) the robustness to visual distractions, 4) the effectiveness of each contributed module, and 5) the transferability to real-world scenarios.

5.1 Experimental Setup

Environments: For training, we employ an augmentable virtual environment called Complex Room. This environment includes a variety of obstacles, distractors, textures, and lighting conditions, providing intricate and diverse scenarios. To further enhance data diversity, we introduce randomness in the size, shape, textures, and lighting of all objects in each episode. The baseline methods of those using online RL are trained in this environment. We also collect the demonstration for offline policy learning, e.g., offline RL or Behavior Cloning.

For testing, we evaluate the agent in five testing environments: 1) Simple Room, serving as a base environment to verify the model’s basic tracking ability. 2) Parking Lot, characterized by obstructed and dim lighting conditions. 3) Urban City, a typical urban street scene with mirror-like reflections on the

ground. 4) Urban Road, similar to Urban City but with lower obstacles. 5) Snow Village, characterized by rough terrain and complex backlight lighting conditions. To evaluate the robustness of the trackers, we introduce varying numbers of distractors into the environment, denoted by the letter “D”. For example, “Parking Lot (4D)” indicates an environment with four distractors. To evaluate the generalization on unseen targets, we add four kinds of animal in Simple Room, including horse, sheep, dog, and pig.

Evaluation: In our evaluation setting, we follow the evaluation setup in previous works [40], which sets the maximum step length in an episode to 500 steps. We define the visible region of the tracker as a 90-degree fan-shaped sector with a radius of 750 cm. A tracker succeeds if it keeps the target in sight for the whole episode (500 steps). The tracker fails if the target continuously leaves this region for more than 50 steps. In each episode, the agent terminates with a success or failure state, and we reset the environment and the agent for the next episode. With the above settings, we employ the following three evaluation metrics to analyze the tracking performance: 1) Accumulated Reward (AR). We assess our model in each environment over 100 episodes and compute the average accumulated reward from these episodes to gauge the tracking performance’s accuracy. 2) Episode Length (EL). We determine the average episode duration over 100 episodes, using termination rules that mirror long-term tracking performance. 3) Success Rate (SR). We calculate the average success rate across 100 evaluation episodes to assess the model’s robustness.

Baselines: We conduct a comprehensive comparison of our proposed approach against eight state-of-the-art models. Specifically, we compare our methods with: 1) DiMP [3]: This model uses a pre-trained video tracker to provide a target bounding box as scene representation and tune a PID controller to control the tracker. 2) SARL [24, 25]: This online RL method encodes RGB input into hidden visual features and trains an end-to-end Conv-LSTM network via reinforcement learning. 3) AD-VAT [42]: This is an asymmetric dueling mechanism that trains the RL-based tracker with a learnable adversarial target to improve the robustness of the end-to-end tracker. 4) AD-VAT+ [41]: The plus version further applies a two-stage training strategy for learning in complex environments with obstacles. 5) RSPT [40]: RSPT is the state-of-the-art RGB-D tracker that additionally uses Depth images to construct a structure-aware representation. It combines an off-the-shelf video tracker, mapping module, and trajectory predictor to construct the representation and train the policy via RL under the AD-VAT mechanism. 6) Behavior Cloning (BC): This is a classic offline policy learning method, which is widely adopted in robotics and autonomous driving. The policy is supervised by expert actions. 7) Cross-modal Teacher-Student Learning (TS) [43]: This is an extension of the BC method by using a pose-based teacher to provide online suggestions for the vision-based student while interacting with environments. Note that the policies for the teacher, target, and distractors are trained in a multi-agent game. 8) Dagger: This is a classic algorithm designed

to iteratively improve policy learning by aggregating datasets of expert actions and the learner’s actions, thereby mitigating the compounding error of the initial imperfect policy. To be fair, we use our proposed text-conditioned segmentation mask as input, the same expert policy as the BC method, and the same training environment. To ensure a fair comparison, we maintain the original parameter settings for all baseline methods.

We also integrate eight configurations for systematic ablation, detailed in Table 3. The first six configurations result from combining various VFM models while omitting specific ablation components. Instance Segmentation Mask (ISM) and RGB data are directly extracted from the game engine (Figure 2). Additionally, we obtained the target bounding box coordinates as visual representations. For each configuration, we use the same offline dataset and adhere to predefined settings for training and evaluation.

Implementation Details: In our experiments, we used a continuous action space for agents. The action space contains two variables: the angular velocity and the linear velocity. Angular velocity varies between $-30^\circ/s$ and $30^\circ/s$, while linear velocity ranges from $-1 m/s$ to $1 m/s$. To train our model, we use an Adam optimizer with a learning rate of $3e-5$ and update the LSTM network every 20 steps. We set the mini-batch size to 32. The whole training process is run on one Nvidia RTX 3090 GPU. For offline policy learning, we collect demonstrations on the Complex Room with 8 distractors. We employ an imperfect expert policy with multi-level action noise (as described in Section 4.2) to generate an offline dataset with 42k steps of data as our training dataset for offline RL training, which we referred to as multi-level imperfect expert dataset. Behavior Cloning (BC) is another offline RL method, which requires perfect expert demonstrations for training. Therefore, we adopt the same procedure but remove all random factors to collect 42k steps of expert dataset for BC.

5.2 Training Time-performance Superiority

The first column of Table 1 indicates the detailed training time of each method. All online baselines (SARL, AD-VAT, AD-VAT+, RSPT, TS, Dagger) take at least 12 hours of training. SARL, AD-VAT+, TS, and Dagger take more than 24 hours to train a comparable tracker. Our proposed method takes 1 hour to outperform all other methods, which shows the advantage of our proposed method in training efficiency. DiMP uses a pre-trained video tracker with a PID controller, which does not need additional training time. However, it needs to manually tune the PID parameters and cannot handle complex environments.

5.3 Generalization to Unseen Environments

We evaluate the tracking performance in 5 different unseen environments and show the results in Table 1. Our method, in conjunction with most baseline approaches, exhibits strong performance within the reference environments (*Simple Room*), except the BC method. We attribute this discrepancy to the inherent

Table 1: Quantitative results compared with baselines in unseen environments. From left to right, the three numbers of each cell represent the Average Accumulated Reward (AR), Average Episode Length (EL), and Success Rate (SR).

Methods/training time	SimpleRoom	Parking Lot	UrbanCity	UrbanRoad	Snow Village	Mean
DiMP / 0 hours	336/ 500 / 1.00	166/327/0.48	239/401/0.66	168/308/0.33	110/301/0.43	204/367/0.58
SARL / 24 hours	368/ 500 / 1.00	92/301/0.22	331/471/0.86	207/378/0.48	203/318/0.31	240/394/0.57
AD-VAT / 12 hours	356/ 500 / 1.00	86/302/0.20	335/484/0.88	246/429/0.60	169/364/0.44	238/416/0.62
AD-VAT+ / 24 hours	373/ 500 / 1.00	267/439/0.60	389 / <u>497</u> / <u>0.94</u>	<u>326</u> /471/0.80	182/365/0.44	307/454/0.76
TS / 26 hours	412 / 500 / 1.00	265/472/ <u>0.89</u>	<u>341</u> /496/ <u>0.94</u>	308/480/0.84	<u>234</u> / <u>424</u> /0.63	<u>312</u> /474/0.86
RSPT / 12 hours	<u>398</u> / 500 / 1.00	314 / <u>480</u> /0.80	<u>341</u> / 500 / 1.00	346 / 500 / 1.00	248 /410/ <u>0.80</u>	329 / <u>478</u> / <u>0.92</u>
Dagger / 24 hours	148/498/0.98	115/446/0.78	133/448/0.78	102/396/0.60	47/154/0.04	109/388/0.64
BC / 1 hours	110/348/0.33	110/350/0.36	182/371/0.47	153/348/0.17	39/279/0.27	119/339/0.32
Ours / 1 hours	374 / 500 / 1.00	<u>274</u> / 484 / 0.92	306 / 500 / 1.00	300/ <u>496</u> / <u>0.96</u>	229/ 471 / 0.87	296/ 490 / 0.95

limitation of the BC method in addressing data shift challenges during offline training. Although Dagger shows comparable long-horizon tracking ability in SimpleRoom, it only achieved a score of 148 in AR, which is lower than other methods. We attribute this to our simple expert policy, which does not provide perfect supervision for learning smooth dynamic control. Consequently, the tracker can keep the target in view, but struggles to control the distance perfectly. These shortcomings are exacerbated in irregular terrains like Snow Village, where the tracker can easily get lost. In unseen environments, our method shows superior long-horizon tracking ability, perfectly following the target in Urban City. It consistently outperformed all baselines except the RSPT method (which uses depth images) in the Urban Road environment. The success rate of our method is only marginally lower by 0.04 in Urban Road. The slight advantage of the RSPT method in Urban Road is due to its depth-based structure-aware motion representation, crucial for navigating low-height obstacles. In environments with complex light conditions and irregular terrains, our approach exceeds RSPT, demonstrating superior performance and robustness to visual distractions.

Table 2: Evaluating the distraction robustness in the environment with distractors. (4D) represent there are 4 distractors in the environment.

Methods	Parking Lot (2D)	UrbanCity (4D)	ComplexRoom (4D)
DiMP	111/271/0.24	170/348/0.32	97/307/0.26
SARL	53/237/0.12	74/221/0.16	22/263/0.15
AD-VAT	43/232/0.13	32/204/0.06	16/223/0.16
AD-VAT+	35/166/0.08	89/245/0.11	35/262/0.18
TS	<u>186</u> / <u>331</u> / <u>0.39</u>	<u>227</u> / <u>381</u> / <u>0.51</u>	<u>250</u> / <u>401</u> / <u>0.54</u>
BC	26/245/0.09	-25/181/0.02	86/336/0.24
Ours	192 / 425 / 0.63	272 / 472 / 0.92	354 / 479 / 0.88

Table 3: The configuration of different settings in state representation. They all are trained via offline RL and tested on the *Complex Room*.

	VFM	Seg	Text	Re-targeting	RNN	AR/EL/SR
Ours	DEVA	✓	✓	✓	✓	306/479/0.92
Ours*	SAM-Track	✓	✓	✓	✓	339/497/0.97
Ours w/o RNN	DEVA	✓	✓	✓		192/413/0.60
Ours w/o re-targeting	DEVA	✓	✓		✓	131/394/0.57
DINOv2	DINOv2	✓			✓	234/434/0.74
E-SAM	E-SAM	✓			✓	-34/253/0.19
ISM	-	✓			✓	8/300/0.23
RGB	-				✓	-63/206/0.03
Bounding box	-				✓	-82/138/0.00

5.4 Generalization to Unseen-category Targets

We evaluate the generalization of our learned policy by tracking different target categories with distractors. We test the trackers in various conditions, including tracking different types of animals in the environment with four distractor animals, denoted as Animals(4D). At each episode, we randomly sample an animal




Fig. 3: **Left:** The exemplars of the visual observation and segmentation mask provided by the vision foundation model in different testing environments. The three snapshots at *Urban Road*, *Urban City*, and *Snow Village* show that the retargeting mechanism can distinguish the target, obstacles, and distractors with different colors in the mask. The sequence at (*Parking Lot*) shows that the VFM can consistently identify the target even when the distractor fully occludes the target. **Right:** The learning curve of different state representations for offline RL validated in *Complex Room* includes the mean and standard deviation of results obtained from training with 3 seeds.

Table 4: Evaluating the generalization on the unseen category of the target in *Simple Room*. We directly adopt the agent on the unseen animals: horse, dog, sheep, and pig.

	Person (4D)	Animals (4D)	Horse	Dog	Sheep	Pig
Ours	285/476/0.90	237/457/0.84	254/500/1.00	231/469/0.90	243/471/0.93	246/472/0.94
Ours w/o RNN	177/454/0.81	207/425/0.77	220/485/0.96	268/465/0.90	251/460/0.92	249/465/0.93
Ours w/o Retargeting	93/343/0.37	27/291/0.25	104/395/0.57	54/319/0.33	32/292/0.30	33/360/0.30
DINOv2	-53/236/0.05	8/237/0.20	137/425/0.71	-22/220/0.07	-31/199/0.03	-2/235/0.09

as the target. We report the performance on Person (4D), where the target and distractors are persons in the SimpleRoom, as the reference. In each setting, the text prompt for VFM is the corresponding category of the target.

The results in Table. 4 show that our proposed method achieved tracking performance similar to that of the reference methods for single-category animals, and even outperformed them in some cases. However, when distractors were present, all three results slightly decreased. This indicates that our model successfully reduced the visual feature gap between categories, but the robustness of dynamic visual distraction across domains can still be improved.

5.5 Robustness to Visual Distraction

We evaluate the robustness of our proposed method in the presence of visual distractions. Such distractors may occlude the tracker or present a texture similar to the target, thereby inducing visual ambiguity. The results in Table 2 show that distractors have a certain degree of impact on the tracker in all environments, but the overall Episode Length of our model remains above 400. In addition, we can see that the success rate of our method significantly outperforms the state-of-the-art distraction-robust method [43] in this setting. This demonstrates that our model has strong anti-distraction abilities and excellent robustness. We also provide an example sequence of our method in case of distraction in the *Parking Lot* environments, as shown in Figure 3.

5.6 Ablation Study

We conduct ablation studies to analyze the effects of each module in our method.

Dataset Size and Noise Level: We assess the impact of the size of the offline dataset by comparing models trained with $20k$ and $80k$ steps. The result illustrates that limited dataset size impedes the performance of policy while incrementing the dataset beyond a threshold yields marginal performance improvement. For behavior cloning, we extend the dataset size to $200k$, $500k$, and $1000k$ steps, observing limited improvement. We also explore the impact of noise level in our approach, which employs a constrained policy to estimate Q values. Pure expert data may lead to Q-function overestimation, and excessive noise can compromise policy quality. Three datasets are collected with varying noise levels, and optimal performance is achieved by noise level 2, cautioning against extreme noise or relying solely on expert data. The implementation details and results are shown in the supplementary material.

Text-conditioned Segmentation Mask: We also analysis the effectiveness of the state representation for offline RL in EVT. We categorize state representations into two groups: those with VFMs and those without. The inference times vary by integrating different models, implementation details can be found in the supplementary material. Our ablation includes DEVA, SAM-Track, E-SAM, and DINOv2, highlighting the benefits of VFMs when combined with text prompts and re-targeting mechanisms. The learning curves are shown in Figure 3, confirming that VFMs with text guidance outperform others, as they effectively reduce background noise and enhance target identification. In particular, DINOv2’s temporal consistency leads to its superior performance over E-SAM. The generalization of our framework is validated across different VFMs (DEVA and SAM-Track), indicating its adaptability to future advancements. Instance segmentation masks or RGB images fall short because of their high-dimensional nature and noise. In contrast, text-conditioned masks bridge the domain gap, improving performance. This study highlights the importance of text-conditioned masks in EVT, enhancing the robustness and effectiveness of our framework.

Recurrent Policy: The RNN in the recurrent policy is also crucial, as the result shown in Table 4. We trained an alternative model without the temporal encoder (w/o RNN). The results indicate a noticeable decline in performance in Person(4D) and Animals(4D), highlighting the effectiveness of temporal information in handling visual distractions.

5.7 Transferring to Real-world Scenarios

We further validate the cross-domain generalization of the learned tracker in real-world videos from the VOT dataset [20] and a mobile robot. We analyze the tracker’s performance by inputting video frames sequentially and comparing predicted actions to human intuition. As shown in Figure 4, the output actions align with our intuition, even when the target is occluded or out-of-view. For instance, in the top sequence, the tracker maintains near-zero angles and high speed during occlusion, then resumes left turns when the deer reappears. We



Fig. 4: The output actions of our agent in two video sequences from VOT Challenge [20]. We separately use “Deer” and “Person” as prompts. The central red point signifies the image’s center. The bottom green rectangle’s width indicates angular velocity control from $-30^\circ/\text{s}$ to $30^\circ/\text{s}$, while the bottom blue rectangle’s height indicates linear velocity control from -1 m/s to 1 m/s .



Fig. 5: The agent is deployed on a real robot, following a human in a complex corridor, navigating corners and hallways, involving pedestrians and a target occlusion.

also deploy our tracker on a mobile robot in an indoor hallway environment, featuring random pedestrian and obstacle distractions. We also present long continuous image sequences to show the robust tracking behavior of the robot controlled by our agent in a real-world corridor with human distractions and static obstacle occlusions, as is shown in Figure 5. The implementation details and more real-world demos can be found in the supplementary material.

6 Conclusion

In this paper, we introduce a novel framework that marries visual foundation models with offline reinforcement learning, leading to an embodied visual tracking agent that is both training-efficient and robust. Our empirical evaluations, conducted across various high-fidelity virtual environments, have demonstrated the framework’s exceptional ability to generalize. It has consistently outperformed existing methods, particularly in scenarios involving previously unseen targets and conditions. These advancements have opened new avenues for exploration and set a benchmark for future work in the domain. Future work will focus on real-world applicability, enhancing adaptability, and minimizing dataset dependence to fully harness the capabilities of embodied agents.

Acknowledgements

This work was supported by the National Science and Technology Major Project (MOST-2022ZD0114900) and China National Post-doctoral Program for Innovative Talents (No. BX2021008).

References

1. Awais, M., Naseer, M., Khan, S., Anwer, R.M., Cholakkal, H., Shah, M., Yang, M.H., Khan, F.S.: Foundational models defining a new era in vision: A survey and outlook. arXiv preprint arXiv:2307.13721 (2023) [4](#)
2. Bajcsy, A., Loquercio, A., Kumar, A., Malik, J.: Learning vision-based pursuit-evasion robot policies. arXiv preprint arXiv:2308.16185 (2023) [3](#)
3. Bhat, G., Danelljan, M., Gool, L.V., Timofte, R.: Learning discriminative model prediction for tracking. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). pp. 6182–6191 (2019) [9](#)
4. Chen, H., Bei, Y., Shen, Q., Xu, Y., Zhou, S., Huang, W., Huang, F., Wang, S., Huang, X.: Macro graph neural networks for online billion-scale recommender systems. In: Proceedings of the ACM on Web Conference 2024. pp. 3598–3608 (2024) [4](#)
5. Chen, H., Xu, Y., Huang, F., Deng, Z., Huang, W., Wang, S., He, P., Li, Z.: Label-aware graph convolutional networks. In: Proceedings of the 29th ACM International Conference on Information & Knowledge Management. p. 1977–1980 (2020) [4](#)
6. Cheng, H.K., Oh, S.W., Price, B., Schwing, A., Lee, J.Y.: Tracking anything with decoupled video segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1316–1326 (2023) [4](#), [5](#), [6](#)
7. Cheng, Y., Li, L., Xu, Y., Li, X., Yang, Z., Wang, W., Yang, Y.: Segment and track anything. arXiv preprint arXiv:2305.06558 (2023) [4](#), [5](#)
8. Ci, H., Liu, M., Pan, X., Zhong, F., Wang, Y.: Proactive multi-camera collaboration for 3d human pose estimation. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=CPIy9TWFYBG> [1](#)
9. Devo, A., Dionigi, A., Costante, G.: Enhancing continuous control of mobile robots for end-to-end visual active tracking. *Robotics and Autonomous Systems* **142**, 103799 (2021) [3](#)
10. Fujimoto, S., Meger, D., Precup, D.: Off-policy deep reinforcement learning without exploration. In: International Conference on Machine Learning. pp. 2052–2062. PMLR (2019) [3](#)
11. Gadre, S.Y., Ehsani, K., Song, S., Mottaghi, R.: Continuous scene representations for embodied ai. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14849–14859 (2022) [4](#)
12. Gao, J., Hu, K., Xu, G., Xu, H.: Can pre-trained text-to-image models generate visual goals for reinforcement learning? *Advances in Neural Information Processing Systems* **36** (2023) [4](#)
13. Haarnoja, T., Zhou, A., Abbeel, P., Levine, S.: Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In: International conference on machine learning. pp. 1861–1870. PMLR (2018) [8](#)

14. Huang, F., Wang, Z., Huang, X., Qian, Y., Li, Z., Chen, H.: Aligning distillation for cold-start item recommendation. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. p. 1147–1157 (2023) **3**
15. Jin, K., Han, X.: Conquering ghosts: Relation learning for information reliability representation and end-to-end robust navigation. arXiv preprint arXiv:2203.09952 (2022) **1**
16. Jin, K., Zhu, H., Gao, R., Wang, J., Wang, H., Yi, H., Shi, C.J.R.: Demrl: Dynamic estimation meta reinforcement learning for path following on unseen unmanned surface vehicle. *Ocean Engineering* **288**, 115958 (2023) **1**
17. Kapelyukh, I., Vosylius, V., Johns, E.: Dall-e-bot: Introducing web-scale diffusion models to robotics. *IEEE Robotics and Automation Letters* (2023) **4**
18. Khandelwal, A., Weihs, L., Mottaghi, R., Kembhavi, A.: Simple but effective: Clip embeddings for embodied ai. In: CVPR. pp. 14829–14838 (2022) **4**
19. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4015–4026 (2023) **4**
20. Kristan, M., Matas, J., Leonardis, A., Vojir, T., Pflugfelder, R., Fernandez, G., Nebehay, G., Porikli, F., Čehovin, L.: A novel performance evaluation methodology for single-target trackers. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38**(11), 2137–2155 (Nov 2016). <https://doi.org/10.1109/TPAMI.2016.2516982> **13, 14**
21. Kumar, A., Zhou, A., Tucker, G., Levine, S.: Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems* **33**, 1179–1191 (2020) **3, 5, 8**
22. Li, J., Xu, J., Zhong, F., Kong, X., Qiao, Y., Wang, Y.: Pose-assisted multi-camera collaboration for active object tracking. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 759–766 (2020) **3**
23. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499 (2023) **4**
24. Luo, W., Sun, P., Zhong, F., Liu, W., Zhang, T., Wang, Y.: End-to-end active object tracking via reinforcement learning. In: International Conference on Machine Learning. pp. 3286–3295. PMLR (2018) **2, 3, 9**
25. Luo, W., Sun, P., Zhong, F., Liu, W., Zhang, T., Wang, Y.: End-to-end active object tracking and its real-world deployment via reinforcement learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **42**(6), 1317–1332 (2020). <https://doi.org/10.1109/TPAMI.2019.2899570> **3, 9**
26. Mousavian, A., Toshev, A., Fišer, M., Košecák, J., Wahid, A., Davidson, J.: Visual representations for semantic target driven navigation. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 8846–8852. IEEE (2019) **4**
27. Puig, X., Undersander, E., Szot, A., Cote, M.D., Yang, T.Y., Partsey, R., Desai, R., Clegg, A.W., Hlavac, M., Min, S.Y., et al.: Habitat 3.0: A co-habitat for humans, avatars and robots. arXiv preprint arXiv:2310.13724 (2023) **3**
28. Qiu, W., Zhong, F., Zhang, Y., Qiao, S., Xiao, Z., Kim, T.S., Wang, Y., Yuille, A.: Unrealcv: Virtual worlds for computer vision. In: Proceedings of the 2017 ACM on Multimedia Conference. pp. 1221–1224 (2017) **6**
29. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from

- natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021) 4
30. Sax, A., Zhang, J.O., Emi, B., Zamir, A., Savarese, S., Guibas, L., Malik, J.: Learning to navigate using mid-level visual priors. In: Conference on Robot Learning. pp. 791–812. PMLR (2020) 4
 31. Shah, D., Bhorkar, A., Leen, H., Kostrikov, I., Rhinehart, N., Levine, S.: Offline reinforcement learning for visual navigation. In: 6th Annual Conference on Robot Learning (2022), https://openreview.net/forum?id=uhIfIEIiWm_ 3
 32. Wang, M., Liu, Y., Su, D., Liao, Y., Shi, L., Xu, J., Miro, J.V.: Accurate and real-time 3-d tracking for the following robots by fusing vision and ultrasonic information. *IEEE/ASME Transactions On Mechatronics* **23**(3), 997–1006 (2018) 1
 33. Xi, M., Zhou, Y., Chen, Z., Zhou, W., Li, H.: Anti-distractor active object tracking in 3d environments. *IEEE Transactions on Circuits and Systems for Video Technology* **32**(6), 3697–3707 (2021) 3
 34. Xiong, Y., Varadarajan, B., Wu, L., Xiang, X., Xiao, F., Zhu, C., Dai, X., Wang, D., Sun, F., Iandola, F., Krishnamoorthi, R., Chandra, V.: Efficientsam: Leveraged masked image pretraining for efficient segment anything (2023) 4
 35. Xu, Y., Chen, H., Wang, Z., Yin, J., Shen, Q., Wang, D., Huang, F., Lai, L., Zhuang, T., Ge, J., Hu, X.: Multi-factor sequential re-ranking with perception-aware diversification. In: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. p. 5327–5337 (2023) 4
 36. Yang, J., Gao, M., Li, Z., Gao, S., Wang, F., Zheng, F.: Track anything: Segment anything meets videos. arXiv preprint arXiv:2304.11968 (2023) 4
 37. Yang, W., Wang, X., Farhadi, A., Gupta, A., Mottaghi, R.: Visual semantic navigation using scene priors. arXiv preprint arXiv:1810.06543 (2018) 4
 38. Yuan, Z., Xue, Z., Yuan, B., Wang, X., Wu, Y., Gao, Y., Xu, H.: Pre-trained image encoder for generalizable visual reinforcement learning. *Advances in Neural Information Processing Systems* **35**, 13022–13037 (2022) 4
 39. Zhang, W., Song, K., Rong, X., Li, Y.: Coarse-to-fine uav target tracking with deep reinforcement learning. *IEEE Transactions on Automation Science and Engineering* **16**(4), 1522–1530 (2018) 3
 40. Zhong, F., Bi, X., Zhang, Y., Zhang, W., Wang, Y.: Rspt: reconstruct surroundings and predict trajectory for generalizable active object tracking. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 3705–3714 (2023) 3, 4, 9
 41. Zhong, F., Sun, P., Luo, W., Yan, T., Wang, Y.: Ad-vat+: An asymmetric dueling mechanism for learning and understanding visual active tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**(5), 1467–1482 (2019) 3, 4, 9
 42. Zhong, F., Sun, P., Luo, W., Yan, T., Wang, Y.: AD-VAT: An asymmetric dueling mechanism for learning visual active tracking. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=HkgYmhR9KX> 2, 3, 9
 43. Zhong, F., Sun, P., Luo, W., Yan, T., Wang, Y.: Towards distraction-robust active visual tracking. In: International Conference on Machine Learning. pp. 12782–12792. PMLR (2021) 3, 7, 9, 12
 44. Zhou, B., Krähenbühl, P., Koltun, V.: Does computer vision matter for action? *Science Robotics* **4**(30), eaaw6661 (2019) 4