
Towards Distraction-Robust Active Visual Tracking

Fangwei Zhong¹ Peng Sun² Wenhan Luo³ Tingyun Yan^{1,4} Yizhou Wang¹

Abstract

In active visual tracking, it is notoriously difficult when distracting objects appear, as distractors often mislead the tracker by occluding the target or bringing a confusing appearance. To address this issue, we propose a mixed cooperative-competitive multi-agent game, where a target and multiple distractors form a collaborative team to play against a tracker and make it fail to follow. Through learning in our game, diverse distracting behaviors of the distractors naturally emerge, thereby exposing the tracker’s weakness, which helps enhance the distraction-robustness of the tracker. For effective learning, we then present a bunch of practical methods, including a reward function for distractors, a cross-modal teacher-student learning strategy, and a recurrent attention mechanism for the tracker. The experimental results show that our tracker performs desired distraction-robust active visual tracking and can be well generalized to unseen environments. We also show that the multi-agent game can be used to adversarially test the robustness of trackers.

1. Introduction

We study Active Visual Tracking (AVT), which aims to follow a target object by actively controlling a mobile robot given visual observations. AVT is a fundamental function for active vision systems and widely demanded in real-world applications, *e.g.*, autonomous driving, household robots, and intelligent surveillance. Here, the agent is required to perform AVT in various scenarios, ranging from a simple room to the wild world. However, the trackers are still vulnerable, when running in an environments with complex situations, *e.g.*, complicated backgrounds, obstacle occlusions, distracting objects.

¹Center on Frontiers of Computing Studies, Dept. of Computer Science, Peking University, Beijing, P.R. China. ²Tencent Robotics X, Shenzhen, P.R. China ³Tencent, Shenzhen, P.R. China ⁴Adv. Inst. of Info. Tech, Peking University, Hangzhou, P.R. China.. Correspondence to: Fangwei Zhong <zfw1226@gmail.com>, Yizhou Wang <yizhou.Wang@pku.edu.cn>.

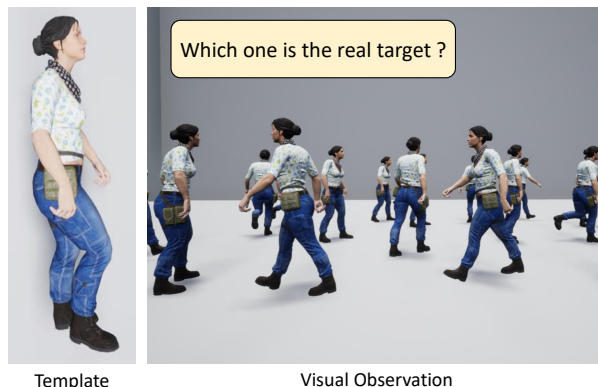


Figure 1. An extreme situation of active visual tracking with distractors. Can you identify which one is the real target?

Among all the challenges, the *distractor*, which can induce confusing visual observation and occlusion, is a considerably prominent difficulty. Distraction emerges frequently in the real world scenario, such as a school where students wear similar uniforms. Considering an extreme case (see Fig. 1), there is a group of people dressed in the same clothes, and you are only given a template image of the target, can you confidently identify the target from the crowd?


The distraction has been primarily studied in the *passive* tracking setting (Kristan et al., 2016) (running on collected video clips) but is rarely considered in the *active* setting (Luo et al., 2020; Zhong et al., 2021). In the *passive* setting, previous researchers (Nam & Han, 2016; Zhu et al., 2018; Bhat et al., 2019) took great efforts on learning a discriminative visual representation in order to identify the target from the crowded background. Yet it is not sufficient for an *active* visual tracker, which should be of not only a suitable state representation but also an optimal control strategy to move the camera to actively avoid distractions, *e.g.*, finding a more suitable viewpoint to seize the target.

To realize such a robust active tracker, we argue that the primary step is to build an interactive environment, where various distraction situations can frequently emergent. A straightforward solution is adding a number of moving objects as distractors in the environment. However, it is non-trivial to model a group of moving objects. Pre-defining object trajectories based on hand-crafted rules seems tempting, but it will lead to overfitting in the sense that the yielded tracker generalizes poorly on unseen trajectories.

Inspired by the recent work on multi-agent learning (Sukhbaatar et al., 2018; Baker et al., 2020), we propose a mixed Cooperative-Competitive Multi-Agent Game to automatically generate distractions. In the game, there are three type of agents: tracker, target and distractor. The tracker tries to always follow the target. The target intends to get rid of the tracker. The target and distractors form a team, which is to make trouble for the tracker. The cooperative-competitive relations among agents are shown in Fig. 2. Indeed, it is already challenging to learn an escape policy for a target (Zhong et al., 2019). As the complicated social interaction among agents, it will be even more difficult to model the behavior of multiple objects to produce distraction situations, which is rarely studied in previous work.

Thus, to ease the multi-agent learning process, we introduce a battery of practical alternatives. First, to mitigate the credit assignments problem among agents, we design a reward structure to explicitly identify the contribution of each distractor with a relative distance factor. Second, we propose a cross-modal teacher-student learning strategy, since directly optimizing the visual policies by Reinforcement Learning (RL) is inefficient. Specifically, we split the training process into two steps. At the first step, we exploit the grounded state to find meta policies for agents. Since the grounded state is clean and low-dimensional, we can easily train RL agents to find a equilibrium of the game, *i.e.*, the target actively explores different directions to escape, the distractors frequently appear in the tracker’s view, and the tracker still closely follows the target. Notably, a multi-agent curriculum is also naturally emergent, *i.e.*, the difficulty of the environment induced by the target-distractors cooperation is steadily increased with the evolution of the meta policies. In the second step, we use the skillful meta tracker (teacher) to supervise the learning of the visual active tracker (student). When the student interacts with the opponents during learning, we replay the multi-agent curriculum by sampling the historical network parameters of the target and distractors. Moreover, a recurrent attention mechanism is employed to enhance the state representation of the visual tracker.

The experiments are conducted in virtual environments with numbers of distractors. We show that our tracker significantly outperforms the state-of-the-art methods in a room with clean backgrounds and a number of moving distractors. The effectiveness of introduced components are validated in ablation study. After that, we demonstrate another use of our multi-agent game, adversarial testing the trackers. While taking adversarial testing, the target and distractors, optimized by RL, can actively find trajectories to mislead the tracker within a very short time period. In the end, we validate that the learned policy is of good generalization in unseen environments. The code and demo videos are available on <https://sites.google.com/view/distraction-robust-avt>.



Cooperative-Competitive Multi-Agent Game

Figure 2. An overview of the Cooperative-Competitive Multi-Agent Game, where the distractors and target cooperate to compete against the tracker.

2. Related Work

Active Visual Tracking. The methods to realize active visual tracking can be simply divided into two categories: two-stage methods and end-to-end methods. Conventionally, the task is accomplished in two cascaded sub-tasks, *i.e.*, the visual object tracking (image \rightarrow bounding box) (Kristan et al., 2016) and the camera control (bounding box \rightarrow control signal). Recently, with the advances of deep reinforcement learning and simulation, the end-to-end methods (image \rightarrow control signal), optimizing the neural network in an end-to-end manner (Mnih et al., 2015), has achieved a great progress in active object tracking (Luo et al., 2018; Li et al., 2020; Zhong et al., 2019). It is firstly introduced in (Luo et al., 2018) to train an end-to-end active tracker in simulator via RL. (Luo et al., 2019) successfully implements the end-to-end tracker in the real-world scenario. AD-VAT(+) (Zhong et al., 2019; 2021) further improves the robustness of the end-to-end tracker with adversarial RL. Following this track, our work is the first paper to study the distractors in AVT with a mixed multi-agent game.

Distractors has attracted attention of researchers in video object tracking (Babenko et al., 2009; Bolme et al., 2010; Kalal et al., 2012; Nam & Han, 2016; Zhu et al., 2018). To address the issue caused by distractors, data augmentation methods like collecting similar but negative samples to train a more powerful classifier is pretty useful (Babenko et al., 2009; Kalal et al., 2012; Hare et al., 2015; Zhu et al., 2018; Deng & Zheng, 2021). Other advanced techniques like hard negative examples mining (Nam & Han, 2016) are also proposed for addressing this. Recently, DiMP (Bhat et al., 2019) develops a discriminative model prediction architecture for tracking, which can be optimized in only a few iterations. Differently, we model the distractors as agents with learning ability, which can actively produce diverse situations to enhance the learning tracker.

Multi-Agent Game. It is not a new concept to apply the multi-agent game to build a robust agent. Roughly speaking,

most of previous methods (Zhong et al., 2019; Florensa et al., 2018; Huang et al., 2017; Mandlekar et al., 2017; Pinto et al., 2017b; Sukhbaatar et al., 2018) focus on modeling a two-agent competition (adversary protagonist) to learn a more robust policy. The closest one to ours is AD-VAT (Zhong et al., 2019), which proposes an asymmetric dueling mechanism (tracker vs. target) for learning a better active tracker. Our multi-agent game can be viewed as an extension of AD-VAT, by adding a number of learnable distractors within the tracker-target competition. In this setting, to compete with the tracker, it necessary to learn collaborative strategies among the distractors and target. However, only a few works (ope; Hausknecht & Stone, 2015; Tampuu et al., 2017; Baker et al., 2020) explored learning policies under a Mixed Multi-Agent Game. In these multi-agent games, the agents usually are homogeneous, *i.e.*, each agent plays an equal role in a team. Differently, in our game, the agents are heterogeneous, including tracker, target, and distractor(s). And the multi-agent competition is also asymmetric, *i.e.*, the active tracker is “lonely”, which has to independently fight against the team formed by the target and distractors.

To find the equilibrium in the multi-agent game, usually, Muti-Agent Reinforcement Learning(MARL) (Rashid et al., 2018; Sunehag et al., 2018; Tampuu et al., 2017; Lowe et al., 2017) are employed. Though effective and successful in some toy examples, the training procedure of these methods is really inefficient and unstable, especially in cases that the agents are fed with high-dimensional raw-pixel observation. Recently, researchers have demonstrated that exploiting the grounded state in simulation can greatly improve the stability and speed of vision-based policy training in single-agent scenario (Wilson & Hermans, 2020; Andrychowicz et al., 2020; Pinto et al., 2017a) by constructing a more compact representation or approximating a more precise value function. Inspired by these, we exploit the grounded state to facilitate multi-agent learning by taking a cross-modal teacher-student learning, which is close to Multi-Agent Imitation Learning (MAIL). MAIL usually rely on a set of collected expert demonstrations (Song et al., 2018; Šošić et al., 2016; Bogert & Doshi, 2014; Lin et al., 2014) or a programmed expert to provide online demonstrations (Le et al., 2017). However, the demonstrations collection and programmed expert designing are usually performed by human. Instead, we adopt a multi-agent game for better cloning the expert behaviour to a vision-based agent. The expert agent is fed with the grounded state and learned by self-play.

3. Multi-Agent Game

Inspired by the AD-VAT (Zhong et al., 2019), we introduce a group of active distractors in the tracker-target competition to induce distractions. We model such a competition as a mixed Cooperative-Competitive Multi-Agent Game, where agents are employed to represent the tracker, target and dis-

tractor, respectively. In this game, the target and distractor(s) constitute a cooperative group to actively find the weakness of the tracker. In contrast, the tracker has to compete against the cooperative group to continuously follow the target. To be specific, each one has its own goal, shown as following:

- *Tracker* chases the target object and keeps a specific relative distance and angle from it.
- *Target* finds a way to get rid of the tracker.
- *Distractor* cooperates with the target and other distractors to help the target escape from the tracker, by inducing confusing visual observation or occlusion.

3.1. Formulation

Formally, we adopt the settings of Multi-Agent Markov Game (Littman, 1994). Let subscript $i \in \{1, 2, 3\}$ be the index of each agent, *i.e.*, $i = 1$, $i = 2$, and $i = 3$ denote the tracker, the target, and the distractors, respectively. Note that, even the number of distractors would be more than one, we use only one agent to represent them. That is because they are homogeneous and share the same policy. The game is governed by the tuple $\langle \mathcal{S}, \mathcal{O}_i, \mathcal{A}_i, \mathcal{R}_i, \mathcal{P} \rangle$, $i = 1, 2, 3$, where $\mathcal{S}, \mathcal{O}_i, \mathcal{A}_i, \mathcal{R}_i, \mathcal{P}$ denote the joint state space, the observation space (agent i), the action space (agent i), the reward function (agent i) and the environment state transition probability, respectively. Let a secondary subscript $t \in \{1, 2, \dots\}$ denote the time step. In the case of visual observation, we have each agent’s observation $o_{i,t} = o_{i,t}(s_t, s_{t-1}, o_{i,t-1})$, where $o_{i,t}, o_{i,t-1} \in \mathcal{O}_i$, $s_t, s_{t-1} \in \mathcal{S}$. Since the visual observation is imperfect, the agents play a partially observable multi-agent game. It reduces to $o_{i,t} = s_t$ in the case of fully observable game, which means that the agent can access the grounded states directly. In the AVT task, the grounded states are the relative poses among all agents, needed by the meta policies. When all the three agents take simultaneous actions $a_{i,t} \in \mathcal{A}_i$, the updated state s_{t+1} is drawn from the environment state transition probability, as $s_{t+1} \sim \mathcal{P}(\cdot | s_t, a_{1,t}, a_{2,t}, a_{3,t})$. Meanwhile, each agent receives an immediate reward $r_{i,t} = r_{i,t}(s_t, a_{i,t})$ respectively. The policy of the agent i , $\pi_i(a_{i,t} | o_{i,t})$, is a distribution over its action $a_{i,t}$ conditioned on its observation $o_{i,t}$. Each policy π_i is to maximize its cumulative discounted reward $\mathbf{E}_{\pi_i} \left[\sum_{t=1}^T \gamma^{t-1} r_{i,t} \right]$, where T denotes the horizontal length of an episode and $r_{i,t}$ is the immediate reward of agent i at time step t . The policy takes as function approximator a neural network with parameter Θ_i , written as $\pi_i(a_{i,t} | o_{i,t}; \Theta_i)$. The cooperation-competition will manifest by the design of reward function $r_{i,t}$, as described in the next subsection.

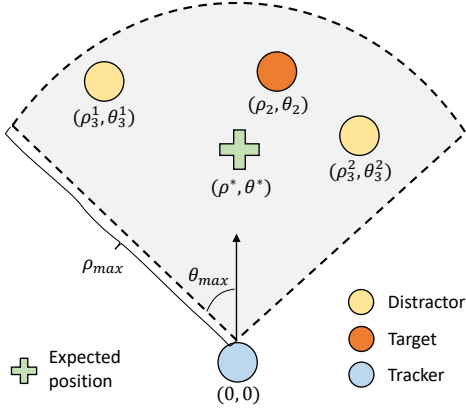


Figure 3. A top-down view of the tracker-centric coordinate system with target (orange), distractors (yellow), and expected target position (green). Tracker (blue) is at the origin of coordinate $(0, 0)$. The gray sector area represents the observable area of the tracker. The arrow in the tracker notes the front of the camera.

3.2. Reward Structure

To avoid the aforementioned intra-team ambiguity, we propose a distraction-aware reward structure, which modifies the cooperative-competitive reward principle to take into account the following intuition: 1) The target and distractors are of a common goal (make the tracker fail). So they need to share the target’s reward to encourage the target-distractor cooperation. 2) Each distractor should have its own rewards, which can measure its contribution to the team. 3) The distractor obviously observed by the tracker will cause distraction in all probability, and the one that out of the tracker’s view can by no means “mislead” the tracker.

We begin with defining a tracker-centric relative distance $d(1, i)$, measuring the geometric relation between the tracker and the other player $i > 1$.

$$d(1, i) = \frac{|\rho_i - \rho^*|}{\rho_{max}} + \frac{|\theta_i - \theta^*|}{\theta_{max}}, \quad (1)$$

where (ρ_i, θ_i) and (ρ^*, θ^*) represent the location of the player $i > 1$ and the expected target in a tracker-centric polar coordinate system, respectively. ρ is the distance from the origin (tracker), and θ is the relative angle to the front of the tracker. See Fig. 3 for an illustration.

With the relative distance, we now give a formal definition of the reward structure as:

$$\begin{aligned} r_1 &= 1 - d(1, 2), \\ r_2 &= -r_1, \\ r_3^j &= r_2 - d(1, j). \end{aligned} \quad (2)$$

Here we omit the timestep subscript t without confusion. The tracker reward is similar to AD-VAT (Zhong et al., 2019), measured by the distance between the target and expected location. The tracker and target play a zero-sum

game, where $r_1 + r_2 = 0$. The distractor is to cooperate with the target by sharing r_2 . Meanwhile, we identify its unique contribution by taking its relative distance $d(1, j)$ as a penalty term in the reward. It is based on an observation that once the tracker is misled by a distractor, the distractor will be regarded as the target and placed at the center of the tracker’s view. Otherwise, the distractor will be penalized by $d(i, j)$ when it is far from the tracker, as its contribution to the gain of r_2 will be marginal. Intuitively, the penalty term $d(1, j)$ can guide the distractors learn to navigate to the tracker’s view, and the bonus from target r_2 can encourage it to cooperate with the target to produce distraction situations to mislead the tracker. In the view of relative distance, the distractor is to minimize $\sum_{t=1}^T d_t(1, j)$ and maximize $\sum_{t=1}^T d_t(1, 2)$, while the tracker is to minimize $\sum_{t=1}^T d_t(1, 2)$. Besides, if a collision is detected to agent i , we penalize the agent with a reward of -1 . When we remove the penalty, the learned distractors would prefer to physically surround and block the tracker, rather than make confusing visual observations to mislead the tracker.

4. Learning Strategy

To efficiently learn policies in the multi-agent game, we introduce a two-step learning strategy to combine the advantages of Reinforcement Learning (RL) and Imitation Learning (IL). First, we train meta policies (using the grounded state as input) via RL in a self-play manner. After that, IL is employed to efficiently impart the knowledge learned by meta policies to the active visual tracker. Using the grounded state can easily find optimal policies (teacher) for each agent first. The teacher can guide the learning of the visual tracker (student) to avoid numerous trial-and-error explorations. Meanwhile, opponent policies emergent in different learning stage are of different level of difficulties, forming a curriculum for the student learning.

4.1. Learning Meta Policies with Grounded State

At the first step, we train meta policies using self-play. Hence, the agent can always play with opponents of an appropriate level, regarded as a natural curriculum (Baker et al., 2020). The meta policies, noted as $\pi_1^*(s_t)$, $\pi_2^*(s_t)$, $\pi_3^*(s_t)$, enjoy privileges to access the grounded state s_t , rather than only visual observations $o_{i,t}$. Even though such grounded states are unavailable in most real-world scenarios, we can easily reach it in the virtual environment. For AVT, the grounded state is the relative poses (position and orientation) among players. We omit the shape and size of the target and distractors, as they are similar during training. Note that the state for agent i will be transformed into the entity-centric coordinate system before feed into the policy network. To be specific, the input of agent i is a sequence about the relative poses to other agents, represented

as $P_{i,1}, P_{i,2}, \dots, P_{i,n}$, where n is the number of the agents and $P_{i,j} = (\rho_{i,j}, \cos(\theta_{i,j}), \sin(\theta_{i,j}), \cos(\phi_{i,j}), \sin(\phi_{i,j}))$. Note that $(\rho_{i,j}, \theta_{i,j}, \phi_{i,j})$ indicates the relative distance, angle, and relative orientation from agent i to agent j . Agent i is at the origin of the coordination. Since the number of the distractors is randomized during either training or testing, the length of the input sequence would be different across each episode. Thus, we adopt the Bidirectional-Gated Recurrent Unit (Bi-GRU) to pin-down a fixed-length feature vector, to enable the network to handle the variable-length distractors. Inspired by the tracker-award model in AD-VAT (Zhong et al., 2019), we also fed the tracker’s action a_1 to the target and distractors, to induce stronger adversaries.

During training, we optimize the meta policies with a modern RL algorithm, *e.g.*, A3C (Mnih et al., 2016). To collect a model pool containing policies at different levels, we save the network parameters the of target and distractor every 50K interactions. During the student learning stage, we can sample the old parameters from the model pool for the target and distractors to reproduce the emergent multi-agent curriculum. Note that we further fine-tune the meta trackers to play against all the opponents in the model pool before going into the next stage. More details about the meta policies can be found in Appendix. A.

4.2. Learning Active Visual Tracking

With the learned meta policies, we seek out a teacher-student learning strategy to efficiently build a distraction-robust active visual tracker in an end-to-end manner, shown as Fig. 4. We apply the meta tracker $\pi_1^*(s_t)$ (teacher) to teach a visual tracker $\pi_1(o_{i,t})$ (student) to track. In the teacher-student learning paradigm, the student needs to clone the teacher’s behavior. Therefore, we dive into the behavioral cloning problem. However, it is infeasible to directly apply supervised learning to learn from the demonstration collected by expert’s behavior. Because the learned policy will inevitably make at least occasional mistakes. However, such a small error may lead the agent to a state which deviates from expert demonstrations. Consequently, The agent will make further mistakes, leading to poor performance. At the end, the student will be of poor generalization to novel scenes.

Thus, we take an interactive training manner as DAGGER (Ross et al., 2011), in which the student takes actions from the learning policy and gets suggestions from the teacher to optimize the policy. To be specific, the training mechanism is composed of two key modules: *Sampler* and *Learner*. In the *Sampler*, we perform the learning policy $\pi_1(o_{i,t})$ to control the tracker to interact with the others. The target and distractors are governed by meta policies $\pi_2^*(s_t)$ and $\pi_3^*(s_t)$ respectively. Meanwhile, the meta tracker $\pi_1^*(s_t)$ provides expert suggestions $a_{1,t}^*$ by monitoring the grounded state. At each step, we sequentially store the visual observa-

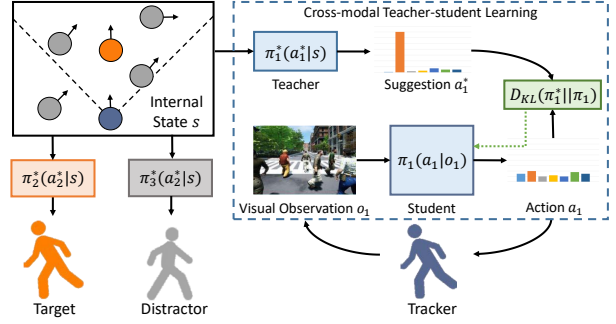


Figure 4. An overview of the cross-modal teacher-student learning strategy. Blue, orange, gray represent tracker, target, distractors, respectively. π_1^* , π_2^* , π_3^* enjoy privileges to acquire the grounded state as input. The tracker adopts the student network (visual tracker) to plays against opponents (target and distractors) to collect useful experiences for learning. We sample parameters from the model pool constructed in the first stage. During training, the student network is optimized by minimizing the KL divergence between the teacher suggestion and the student output.

tion and the expert suggestions $(o_{1,t}, a_{1,t}^*)$ in a buffer B . To make diverse multi-agent environments, we random sample parameters from the model pool for $\pi_2^*(s_t)$ and $\pi_3^*(s_t)$. The model pool is constructed during the first stage, containing meta policies at different levels. So we easily reproduce the multi-agent curriculum emergent in the first stage. We also demonstrate the importance of the multi-agent curriculum in the ablation analysis.

In parallel, the *Learner* samples a batch of sequences from the buffer B and optimizes the student network in a supervised learning manner. The objective function of the learner is to minimize the relative entropy (Kullback-Leibler divergence) of the action distribution between student and teacher, computed as:

$$\mathcal{L}_{KL} = \sum_{n=1}^N \sum_{t=1}^T D_{KL}(a_{1,t}^* || \pi(o_t)), \quad (3)$$

where N is the number of trajectories in the sampled batch, T is the length of one trajectory. In practice, multiple samplers and one learner work asynchronously, significantly reducing the time needed to obtain satisfactory performance.

Moreover, we employ a recurrent attention mechanism in the end-to-end tracker network (Zhong et al., 2019) to learn a representation which is consistent in spatial and temporal. We argue that a spatial-temporal representation is needed for the active visual tracking, especially in the case of distraction appearing. Specially, we use the ConvLSTM (Xingjian et al., 2015) to encode an attention map, which is multiplied by the feature map extracted a target-aware feature from the CNN encoder. See Appendix.B for more details.

5. Experimental Setup

In this section, we introduce the environments, baselines, and evaluation metrics used in our experiments.

Environments. Similar to previous works (Luo et al., 2020; Zhong et al., 2019), the experiments are conducted on UnrealCV environments (Qiu et al., 2017). We extend the two-agent environments used in AD-VAT (Zhong et al., 2019) to study the multi-agent ($n > 2$) game. Similar to AD-VAT, the action space is discrete with seven actions, *move-forward/backward*, *turn-left/right*, *move-forward-and-turn-left/right*, and *no-op*. The observation for the visual tracker is the color image in its first-person view. The primary difference is that we add a number of controllable distractors in the environment, shown as Fig. 5. Both target and distractors are controlled by the scripted navigator, which temporally sample a free space in the map and navigate to it with a random set velocity. Note that we name the environments that use the scripted navigator to control the target and x distractors as *Nav- x* . If they are governed by the meta policies, we mark the environment as *Meta- x* . Besides, we enable agents to access the poses of players, which is needed by the meta policies. Two realistic scenarios (*Urban City* and *Parking Lot*) are used to verify the generalization of our tracker in other unseen realistic environments with considerable complexity. In *Urban City*, there are five unseen characters are placed, and the appearance of each is randomly sampled from four candidates. So it is potential to see that two characters dressed the same in an environment. In *Parking Lot*, all of the target and distractors are of the same appearance. Under this setting, it would be difficult for the tracker to identify the target from distractions.

Evaluation Metrics. We employ the metrics of Accumulated Reward (AR), Episode Length (EL), Success Rate (SR) for our evaluation. Among those metrics, AR is the recommended primary measure of tracking performance, as it considers both precision and robustness. The other metrics are also reported as auxiliary measures. Specifically, AR is affected by the immediate reward and the episode length. Immediate reward measures the goodness of tracking at each step. EL roughly measures the duration of good tracking, as the episode is terminated when the target is lost for continuous 5 seconds or it reaches the max episode length. SR is employed in this work to better evaluate the robustness, which counts the rate of successful tracking episodes after running 100 testing episodes. An episode is marked as success only if the tracker continuously follows the target till the end of the episode (reaching the max episode length).

Baselines. We compare our method with a number of state-of-the-art methods and their variants, including the two-stage and end-to-end trackers. First, We develop conventional two-stage active tracking methods by combining passive trackers with a PID-like controller. As for the passive



Figure 5. The snapshots of tracker’s visual observation in *Simple Room*. The right is the augmented training environment. The target is pointed out with a bounding box.

trackers, we directly use three off-the-shelf models (DaSi-amRPN (Zhu et al., 2018), ATOM (Danelljan et al., 2019), DiMP (Bhat et al., 2019)) without additional training in our environment. Notably, both DiMP and ATOM can be optimized on the fly to adapt to a novel domain. So they can generalize well to our virtual environments and achieve strong performance in the no-distractor environments, *e.g.*, DiMP tracker achieves 1.0 SR in *Simple Room (Nav-0)*. Second, two recent end-to-end methods (SARL (Luo et al., 2020), AD-VAT (Zhong et al., 2019)) are reproduced in our environment to compare. We also extend them by adding two random walking distractors in the training environment, noted as SARL+ and AD-VAT+.

All end-to-end methods are trained in *Simple Room* with environment augmentation (Luo et al., 2020). After the learning converged, we choose the model that achieves the best performance in the environment for further evaluation. Considering the random factors, we report the average results after running 100 episodes. More implementation details are introduced in Appendix.C.

6. Results

We first demonstrate the evolution of the meta policies while learning in our game. Then, we report the testing results in *Simple Room* with different numbers of distractors. After that, we conduct an ablation study to verify the contribution of each component in our method. We also adversarially test the trackers in our game. Moreover, we evaluate the transferability of the tracker in photo-realistic environments.

6.1. The Evolution of the Meta Policies

While learning to play the multi-agent game, the multi-agent curriculum automatically emerges. To demonstrate it, we evaluate the skill-level of the adversaries (target+distractors) at different learning stages from two aspects: the frequency of the distractor appearing in the tracker’s view and the success rate (SR) of the off-the-shelf trackers (DiMP and ATOM). To do it, we collect seven meta policies after agents take 0, 0.4M, 0.7M, 1M, 1.3M, 1.65M, 2M interactions, respectively. We then make the visual tracker (ATOM and DiMP) to play with each collected adversaries

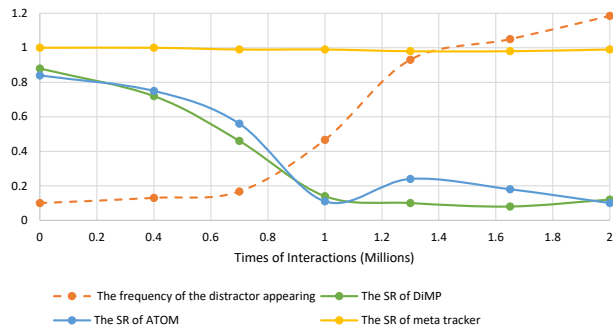


Figure 6. The evolution of target-distractor cooperation in the multi-agent game.

(one target and two distractors) in *Simple Room*, and count the success rate of each tracker in 100 episodes. We also let the converged meta tracker (at $2M$) follow the target-distractors group from different stages, respectively. And we report the success rate of the meta tracker and the frequency of the distractor appearing in the tracker’s view, shown as Fig. 6. We can see that the frequency of the distractor appearing is increased during the multi-agent learning. Meanwhile, the success rate of the visual tracker is decreased. This evidence shows that the complexity of the multi-agent environment is steadily increased with the development of the target-distractor cooperation. Meanwhile, we also notice that the learned meta tracker can robustly follow the target ($SR \geq 0.98$), even when the distractors frequently appearing in its view. This motivates us to take the meta tracker as a teacher to guide the learning of the active visual tracker.

6.2. Evaluating with Scripted Distractors

We analyze the distraction robustness of the tracker in *Simple Room* with scripted target and distractors. This environment is relatively simple to most real-world scenarios as the background is plain and no obstacles is placed. So most trackers can precisely follow the target when there is no presence of distraction. Hence, we can explicitly analyze the distraction robustness by observing how the tracker’s performance is changed with the increasing number of distractors, shown as Fig. 7. Note that we normalize the reward by the average score achieved by the meta tracker, which is regarded as the best performance that the tracker could reach in each configuration. We observed that the learned meta tracker is strong enough to handle different cases, and hardly lost in the environment.

We can see that most methods (except DaSiamRPN) are competitive when there is no distractor. The low score of DaSiamRPN is mainly due to the inaccuracy of the predicted bounding box, which further leads to the tracker’s failure in keeping a certain distance from the target. With the increasing number of distractors, the gap between our tracker and baselines gradually broaden. For example, in




Figure 7. Evaluating the distraction-robustness of the trackers by increasing the number of random distractors in *Simple Room*

the four distractors room, the normalized reward achieved by our tracker is two times the ATOM tracker, *i.e.*, 0.61 vs. 0.25. For the two-stage methods, in most simple cases, DiMIP is a little better than ATOM, thanks to the discriminative model prediction. However, it remains to get lost easily when the distractor occludes the target. We argue that there are two reasons leading to the poor performance of the passive tracker when distractions appearing: 1) the target representation is without long-term temporal context. 2) the tracker lacks a mechanism to predict the state of an invisible target. Thus, once the target is disappeared (occluded by distractors or out of the view), the tracker will regard the observed distractor as a target. If it follows the false target to go, the true target will hardly appear in its view again.

For the end-to-end methods, it seems much weaker than the two-stage methods in the distraction robustness, especially when playing against many distractors. By visualizing the tracking process, we find that AD-VAT tends to follow the moving object in the view but unable to identify which is the true target. So it is frequently misled by the moving objects around the target. Besides, the curves of SARL+ and AD-VAT+ are very close to the original version (SARL and AD-VAT). However, without the target-distractor cooperation, the distraction situation appears at a low frequency in the plus versions. Thus, the learned trackers are still vulnerable to the distractors, and the improvements they achieved are marginal. This indicates that it is useless to simply augment the training environment with random moving distractors.

Our tracker significantly outperforms others in all the cases of distractors. This evidence shows that our proposed method is of great potential for realizing robust active visual tracking in a complex environment. However, the performance gap between our model and the teacher model ($1 - ours$) indicates that there is still room for improvement. By visualizing the test sequences of our model we find that it mainly fails in extreme tough cases where the tracker is surrounded by some distractors that totally occlude the target or block the way to track. More vivid examples can be found in the demo video.

Table 1. Ablative analysis of the visual policy learning method in *Simple Room*. The best results are shown in bold.

Methods	Nav-4			Meta-2		
	AR	EL	SR	AR	EL	SR
Ours	250	401	0.54	141	396	0.44
w/o multi-agent curriculum	232	394	0.53	-23	283	0.08
w/o teacher-student learning	76	290	0.22	79	340	0.4
w/o recurrent attention	128	193	0.27	75	331	0.32

6.3. Ablation Study

We conduct an ablation study to better understand the contribution of each introduced component in our learning method. 1) To evaluate the effectiveness of the multi-agent curriculum, we use the scripted navigator to control the target and distractors when taking the teacher-student learning, instead of replaying the policies collected in the model pool. We find that the learned tracker obtains comparable results to ours in the Nav-4, but there is an obvious gap in Meta-2, where the target and 2 distractors are controlled by the adversarial meta policies. This shows that the tracker over-fits specific moving pattern of the target, but does not learn the essence of active visual tracking. 2) For teacher-student learning, we directly optimize the visual tracking network by A3C, instead of using the suggestions from the meta tracker. Notably, such method can also be regarded as a method that augments the environment by multi-agent curriculum for SARL method. So, by comparing its result with SARL, we can also recognize the value of the multi-agent curriculum on improving the distraction robustness. For the recurrent attention, we compare it with the previous ConvLSTM network introduced in (Zhong et al., 2019). We can see that the recurrent attention mechanism can significantly improve the performance of the tracker in both settings.

6.4. Adversarial Testing

Beyond training a tracker, our multi-agent game can also be used as a test bed to further benchmark the distraction robustness of the active trackers. In adversarial testing, the target collaborates with distractors to actively find adversarial trajectories that fail the tracker to follow. Such an adversarial testing is necessary for AVT. Because the trajectories generated by rule-based moving objects designed for evaluation can never cover all of the possible cases, *i.e.* the trajectories of objects can be arbitrary and have infinitely possible patterns. Moreover, it can also help us discover and understand the weakness of the learned tracker, thus facilitating further development.

We conduct the adversarial testing by training the adversaries to find model-specific adversarial trajectories for each tracker. In this stage, the network of target and distractors are initialized with parameters from the meta policies. For a fair comparison, We iterate the adversaries in 100K inter-




Figure 8. The reward curves of four trackers during the adversarial testing, running with three random seeds. Better viewed in color.

action samples. The model of the tracker is frozen during the adversarial testing. The adversarial testing is conducted in *Simple Room* with 2 distractors. The curves are the average of three runs with different random seeds, and the shaded areas are the standard errors of the mean. Fig. 8 plots the reward of four trackers during the adversarial testing.

We find that most of the trackers are vulnerable to adversaries, resulting in a fast descending of the model’s reward. The rewards of all the methods drop during the testing, showing that the adversaries are learning a more competitive behaviour. The reward of our tracker leads the baseline methods most of the time. In the end, DiMP and ATOM are struggling in the adversarial case, getting very low rewards, ranging from -60 to -100 .

Besides, we also observe an interesting but difficult case. The target rotates at a location and the distractors move around the target and occlude the target; After a while, the distractor goes away, and the two-stage trackers will follow the distractor instead of the target. The demo sequences are available in the demo video. The adversarial testing provides a new evidence to the robustness of our tracker. It also reflects the effectiveness of our method in learning target-distractor collaboration.

6.5. Transferring to Unseen Environments

To show the potential of our model in realistic scenarios, we validate the transferability of the learned model in two photo-realistic environments, which are distinct from the training environment.

As the complexity of the environment increases, performance of these models is downgraded comparing to the results in *Simple Room*, shown as Fig. 9. Even so, our tracker still significantly outperforms others, showing the stronger transferability of our model. In particular, in *Parking Lot* where the target and distractor have the same appearance, the tracker must be able to consider the spatial-temporal




Figure 9. Evaluating generalization of the tracker on two unseen environments (*Urban City* and *Parking Lot*).



Figure 10. Two exemplar sequences of our tracker running on the *Urban City* (top) and *Parking Lot* (bottom). For better understanding, we point out the target object with a red line bounding box in the first frame. More examples are available in the demo video.

consistency to identify the target. Correspondingly, trackers which mainly rely on the difference in appearance is not capable of perceiving the consistency, there they should not perform well. In contrast, our tracker performs well in such cases, from which we can infer that our tracker is able to learn spatial-temporal consistent representation that can be very useful when transferring to other environments. Two typical sequences are shown in Fig. 10.

7. Conclusion and Discussion

Distracting objects are notorious for degrading the tracking performance. This paper offers a novel perspective on how to effectively train a distraction-robust active visual tracker, which is a problem setting that has barely been addressed in previous work. We propose a novel multi-agent game for learning and testing. Several practical techniques are introduced to further improve the learning efficiency, including designing reward function, two-stage teacher-student learning strategy, recurrent attention mechanism *etc.*. Empirical results on 3D environments verified that the learned tracker is more robust than baselines in the presence of distractors.

Considering a clear gap between our trackers and the ideal one (ours vs teacher model), there are many interesting future directions to be explored beyond our work. For example, we can further explore a more suitable deep neural network for the visual tracker. For real-world deployment, it is also necessary to seek an unsupervised or self-supervised domain adaption method (Hansen et al., 2021) to improve

the adaptive ability of the tracker on novel scenarios. Besides, it is also feasible to extend our game on other settings or tasks, such as multi-camera object tracking (Li et al., 2020), target coverage problem (Xu et al., 2020), and moving object grasping (Fang et al., 2019).

Acknowledgements

Fangwei Zhong, Tingyun Yan and Yizhou Wang were supported in part by the following grants: MOST-2018AAA0102004, NSFC-61625201, NSFC-62061136001, the NSFC/DFG Collaborative Research Centre SFB/TRR169 “Crossmodal Learning” II, Tencent AI Lab Rhino-Bird Focused Research Program (JR201913), Qualcomm University Collaborative Research Program.

References

- Openai five. <https://blog.openai.com/openai-five/>. Accessed August 30, 2018.
- Andrychowicz, O. M., Baker, B., Chociej, M., Jozefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.
- Babenko, B., Yang, M.-H., and Belongie, S. Visual tracking with online multiple instance learning. In *The IEEE Conference on Computer Vision and Pattern Recognition*, pp. 983–990, 2009.
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. Emergent tool use from multi-agent autotutorials. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkxpxJBkWS>.
- Bhat, G., Danelljan, M., Gool, L. V., and Timofte, R. Learning discriminative model prediction for tracking. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 6182–6191, 2019.
- Bogert, K. and Doshi, P. Multi-robot inverse reinforcement learning under occlusion with interactions. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 173–180, 2014.
- Bolme, D. S., Beveridge, J. R., Draper, B. A., and Lui, Y. M. Visual object tracking using adaptive correlation filters. In *The IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2544–2550, 2010.
- Danelljan, M., Bhat, G., Khan, F. S., and Felsberg, M. Atom: Accurate tracking by overlap maximization. In *Proceed-*

- ings of the *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4660–4669, 2019.
- Deng, W. and Zheng, L. Are labels always necessary for classifier accuracy evaluation? In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2021.
- Fang, M., Zhou, C., Shi, B., Gong, B., Xu, J., and Zhang, T. DHER: Hindsight experience replay for dynamic goals. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Byf5-30qFX>.
- Florensa, C., Held, D., Geng, X., and Abbeel, P. Automatic goal generation for reinforcement learning agents. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pp. 1515–1528, 2018.
- Hansen, N., Jangir, R., Sun, Y., Alenyà, G., Abbeel, P., Efros, A. A., Pinto, L., and Wang, X. Self-supervised policy adaptation during deployment. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=o_V-MjyyGV_.
- Hare, S., Golodetz, S., Saffari, A., Vineet, V., Cheng, M.-M., Hicks, S. L., and Torr, P. H. Struck: Structured output tracking with kernels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(10):2096–2109, 2015.
- Hausknecht, M. and Stone, P. Deep reinforcement learning in parameterized action space. *arXiv preprint arXiv:1511.04143*, 2015.
- Huang, S., Papernot, N., Goodfellow, I., Duan, Y., and Abbeel, P. Adversarial attacks on neural network policies. *arXiv preprint arXiv:1702.02284*, 2017.
- Kalal, Z., Mikolajczyk, K., and Matas, J. Tracking-learning-detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(7):1409–1422, 2012.
- Kristan, M., Matas, J., Leonardis, A., Vojir, T., Pflugfelder, R., Fernandez, G., Nebehay, G., Porikli, F., and 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.
- Le, H. M., Yue, Y., Carr, P., and Lucey, P. Coordinated multi-agent imitation learning. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 1995–2003, 2017.
- Li, J., Xu, J., Zhong, F., Kong, X., Qiao, Y., and Wang, Y. Pose-assisted multi-camera collaboration for active object tracking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 759–766, 2020.
- Lin, X., Beling, P. A., and Cogill, R. Multi-agent inverse reinforcement learning for zero-sum games. *arXiv preprint arXiv:1403.6508*, 2014.
- Littman, M. L. Markov games as a framework for multi-agent reinforcement learning. In *Machine Learning Proceedings 1994*, pp. 157–163. Elsevier, 1994.
- Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, O. P., and Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Advances in Neural Information Processing Systems*, pp. 6379–6390, 2017.
- Luo, W., Sun, P., Zhong, F., Liu, W., Zhang, T., and Wang, Y. End-to-end active object tracking via reinforcement learning. In *International Conference on Machine Learning*, pp. 3286–3295, 2018.
- Luo, W., Sun, P., Zhong, F., Liu, W., Zhang, T., and 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, 2019.
- Luo, W., Sun, P., Zhong, F., Liu, W., Zhang, T., and 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.
- Mandlekar, A., Zhu, Y., Garg, A., Fei-Fei, L., and Savarese, S. Adversarially robust policy learning: Active construction of physically-plausible perturbations. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3932–3939. IEEE, 2017.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, 2015.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pp. 1928–1937, 2016.
- Nam, H. and Han, B. Learning multi-domain convolutional neural networks for visual tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4293–4302, 2016.

- Pinto, L., Andrychowicz, M., Welinder, P., Zaremba, W., and Abbeel, P. Asymmetric actor critic for image-based robot learning. *arXiv preprint arXiv:1710.06542*, 2017a.
- Pinto, L., Davidson, J., Sukthankar, R., and Gupta, A. Robust adversarial reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 2817–2826. JMLR.org, 2017b.
- Qiu, W., Zhong, F., Zhang, Y., Qiao, S., Xiao, Z., Kim, T. S., Wang, Y., and Yuille, A. Unrealcv: Virtual worlds for computer vision. In *Proceedings of the 2017 ACM on Multimedia Conference*, pp. 1221–1224, 2017.
- Rashid, T., Samvelyan, M., Schroeder, C., Farquhar, G., Foerster, J., and Whiteson, S. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 4295–4304. PMLR, 2018.
- Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 627–635, 2011.
- Song, J., Ren, H., Sadigh, D., and Ermon, S. Multi-agent generative adversarial imitation learning. In *Advances in Neural Information Processing Systems*, pp. 7461–7472, 2018.
- Šošić, A., KhudaBukhsh, W. R., Zoubir, A. M., and Koepl, H. Inverse reinforcement learning in swarm systems. *arXiv preprint arXiv:1602.05450*, 2016.
- Sukhbaatar, S., Lin, Z., Kostrikov, I., Synnaeve, G., Szlam, A., and Fergus, R. Intrinsic motivation and automatic curricula via asymmetric self-play. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SkT5Yg-RZ>.
- Sunehag, P., Lever, G., Grusl, A., Czarnecki, W. M., Zambaldi, V., Jaderberg, M., Lanctot, M., Sonnerat, N., Leibo, J. Z., Tuyls, K., et al. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 2085–2087, 2018.
- Tampuu, A., Matiisen, T., Kodelja, D., Kuzovkin, I., Korjus, K., Aru, J., Aru, J., and Vicente, R. Multiagent cooperation and competition with deep reinforcement learning. *PLoS one*, 12(4):e0172395, 2017.
- Wilson, M. and Hermans, T. Learning to manipulate object collections using grounded state representations. In *Conference on Robot Learning*, pp. 490–502, 2020.
- Xingjian, S., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., and Woo, W.-c. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in Neural Information Processing Systems*, pp. 802–810, 2015.
- Xu, J., Zhong, F., and Wang, Y. Learning multi-agent coordination for enhancing target coverage in directional sensor networks. In *Advances in Neural Information Processing Systems*, volume 33, pp. 10053–10064, 2020.
- Zhong, F., Sun, P., Luo, W., Yan, T., and Wang, Y. AD-VAT: An asymmetric dueling mechanism for learning visual active tracking. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=HkgYmhR9KX>.
- Zhong, F., Sun, P., Luo, W., Yan, T., and Wang, Y. Advat+: An asymmetric dueling mechanism for learning and understanding visual active tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(5):1467–1482, 2021.
- Zhu, Z., Wang, Q., Li, B., Wu, W., Yan, J., and Hu, W. Distractor-aware siamese networks for visual object tracking. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 101–117, 2018.