Active visual tracking in multi-agent scenarios

Yiming Wang and Andrea Cavallaro
Centre for Intelligent Sensing, Queen Mary University of London
{yiming.wang, a.cavallaro}@qmul.ac.uk

Abstract

We propose an active visual tracker with collision avoidance for camera-equipped robots in dense multi-agent scenarios. The objective of each tracking agent (robot) is to maintain visual fixation on its moving target while updating its velocity to avoid other agents. However, when multiple robots are present or targets intensively intersect each other, robots may have no accessible collision-avoiding paths. We address this problem with an adaptive mechanism that sets the pair-wise responsibilities to increase the total accessible collision-avoiding controls. The final collision-avoiding control accounts for motion smoothness and view performance, i.e. maintaining the target centered in the field of view and at a certain size. We validate the proposed approach under different target-intersecting scenarios and compare it with the Optimal Reciprocal Collision Avoidance and the Reciprocal Velocity Obstacle methods.

1. Introduction

Robots that actively track a person can offer assistance in shops, airports and museums. Active tracking with a camera aims to maintain the target at a desired position and size on the image plane [12] by either optimization [19] or various feedback controllers [8, 12]. The motion of an active visual tracker is constrained by its limited field of view and the target dynamics. When multiple robots coexist in a scene and track their own target, collisions are likely to occur when the paths of the targets intersect each other. For multi-robot collision avoidance, as robots react to the actions of other robots, reciprocal avoidance strategies are essential to avoid oscillations.

Methods exist that apply human strategies for safe robot navigation, where the strategy can be learned from extracted trajectories [10] or based on cognitive studies [11]. A popular method for multi-robot collision avoidance is Optimal Reciprocal Collision Avoidance (ORCA) [17]. ORCA is a velocity-based method that guarantees collision-free motion in a densely-packed environment. Each robot independently computes the closest collision-avoiding velocity to its preferred velocity (i.e. the velocity the robot would move at if there were no other robots in its way). Assuming that each robot can sense the state of neighboring robots and infer their preferred velocities, ORCA does not require communication among robots.

When targets are intensively intersecting, robots may face the empty-set problem (i.e. without accessible collision-free velocities [17]). The original ORCA adjusts the pair-wise velocity constraints in order to have at least one accessible velocity in the case of an empty-set occurrence [17]. However, by doing so the robot motion is completely dependent on neighboring robots. Instead, we prevent the occurrences of the empty-set problem by adapting the pair-wise responsibilities. The pair-wise responsibility can assign the right of way to the agent with a constant higher priority [6], while we set the pair-wise responsibility at run time to reduce the responsibility in avoiding other robots for a robot with a smaller set of accessible collision-avoiding velocities.

In this paper, we propose an adaptive ORCA (A-ORCA) method that addresses the empty-set problem under challenging target-intersecting scenarios. The proposed method adapts the pair-wise responsibility with the objective of increasing the accessible collision-avoiding velocities of a pair of robots in a fair way. The final collision-avoiding control is achieved via an optimal controller that accounts for both view performance and motion smoothness, and satisfies the velocity constraints derived with A-ORCA. We compare the proposed method with ORCA [17] and RVO [16] under different target-intersecting scenarios.

2. Related work

ORCA is based on the concept of Velocity Obstacle (VO). VO was originally proposed for a robot to avoid moving obstacles with known paths by defining the set of velocities that can lead to collisions within a time horizon [7]. In multi-robot scenarios, each robot is a *reactive* moving obstacle for other robots. Therefore oscillations are likely to occur if a robot avoids other robots without accounting for their reactive nature. Reciprocal Velocity Obstacle (RVO) avoids reactive robots by using the average of the preferred

Table 1. Multi-agent collision avoidance methods based on Optimal Reciprocal Collision Avoidance. KEY – Ref: reference; AVT: Active Visual Tracking; Nav: Navigation; Holo: Holonomic; Nonholo: Nonholonomic; Prop: Proposed method.

Ref	Application		Agent shape		Agent kinematics		Equal
	AVT	Nav	Circular	Other	Holo	Nonholo	responsibility
[2]		✓	✓			✓	✓
[3]		✓		Cylindrical		✓	✓
[4]		✓	✓			✓	✓
[5]		✓		Elliptical	✓		✓
[6]		✓	✓		✓		Priority-based
[15]		✓	✓			✓	✓
[16]		✓	✓		✓		✓
[17]		✓	✓		✓		✓
Prop	✓		✓			✓	Adaptive

velocities together with VO, in order to obtain a set of reciprocal collision-avoiding velocities [16]. Based on RVO, ORCA further defines the set of velocities that are both reciprocally collision-avoiding and close to the preferred velocity [17]. Each robot shares equal responsibility to avoid another robot. The original ORCA considers disk-shaped robots with a holonomic model constrained by a speed limit (a circular boundary in the velocity space). Derivative works extend ORCA in terms of robot shapes [3, 5] and robotic kinematics [2, 3, 4, 15].

ORCA-related works apply to robots navigating to fixed goal positions [2, 4, 5, 15, 16, 17]. In the field of active visual tracking, the main focus is the control for visual fixation [8, 12] or for minimizing tracking uncertainty [19]. While there is work for multiple robots tracking one target, collision avoidance is achieved via maintaining a predefined formation [13].

To the best of our knowledge, no works have addressed multi-robot collision avoidance among robots that are independently conducting active visual tracking. Collision avoidance is more challenging during view maintenance as the paths of the targets may frequently intersect with each other.

Methods for multi-robot collision avoidance with ORCA [2, 4, 5, 15, 16, 17] are compared in Table 1.

3. Collision-free active visual tracking

3.1. Problem definition

Let $\{c_1, c_2, ..., c_i, ..., c_M\}$ be M disk-shaped robots that aim to track M point targets. Each robot is actively tracking one target. Let i be the index for robots, n the index for targets and t the time step. Each robot has radius r, position $\mathbf{p}_i(t)$ and velocity $\mathbf{v}_i(t)$ in the global coordinate system. Let the distance between c_i and its target be $d_{in}(t) = \|\mathbf{p}_n(t) - \mathbf{p}_i(t)\|$. Let $\delta_{in}(t) \in (-\pi,\ \pi]$ be the angle between the robot heading direction and its target (see Fig. 1).

Each robot is equipped with a camera whose field of view (FoV) is a sector and whose orientation is the same as the robot heading direction. Each c_i aims to fixate its target by maintaining it at a certain distance away in the camera heading direction while avoiding collisions with any other robot c_j, i.e. $d_{ij}(t) > 2r, \forall j,\ j \neq i$.

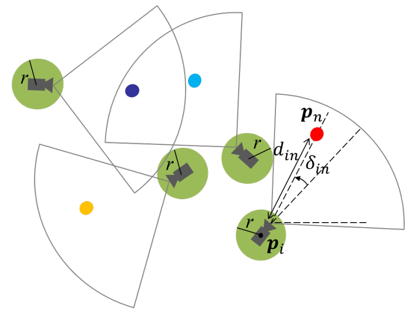

Figure 1. Top-view for active visual tracking with four robots of radius r. Point-targets are indicated with different colors that correspond to their IDs. Each robot has a camera with a sector-shaped field of view facing the same direction as the robot heading direction. The robot located at $\mathbf{p}_i$ is tracking its target at $\mathbf{p}_n$. d_{in} is the distance from the robot to the target and δ_{in} is the angle from the robot heading direction to the target.

We assume that the assignment between robots and targets is given and ambient cameras localize targets and robots. These positions are communicated to the robots [4, 14, 18]. At initialization, robots exchange (or are given) their preferred velocities. Using the positions of nearby robots and the target to follow, a robot calculates at each time step, t, its preferred velocity for active tracking in order to maintain the target at the desired position in the FoV, i.e. the FoV center [18], in a smooth manner.

Using the positions and preferred velocities of neighboring robots, each c_i derives the set of velocities that guarantees no collisions in τ time steps and computes the final control within that set. A short time horizon τ leads to late aversion of avoidance, which can result in empty-set cases as the closeness among robots leaves no space to move; whereas a large τ averts early avoidance that leaves more space to move, but imposes restrictive velocity constraints which can also cause empty sets [4].

3.2. The controller

We use a car-like kinematic model [2]. The robotic control $\mathbf{a}_i(t)$ consists of speed $v_i(t)$ and steering angle $\omega_i(t)$, and is bounded as $|v_i(t)| \leqslant v_{\max}$ and $|\omega_i(t)| \leqslant \omega_{\max}$. We compute the robotic control $\mathbf{a}_i(t)$ by minimizing two cost functions that encode the objectives of maintaining the target at the FoV center and motion smoothness.

Let $\rho^d_{in}(t+1) = \frac{2\left|d_{in}(t+1) - \frac{1}{2}r_v\right|}{r_v}$ and $\rho^\delta_{in}(t+1) = \frac{2\delta_{in}(t+1)}{\phi}$ be the ratio of the distance and angle deviation to the FoV center at $t+1$ with respect to half view range $\frac{r_v}{2}$ and half view angle $\frac{\phi}{2}$. We estimate $d_{in}(t+1)$ and $\delta_{in}(t+1)$

based on the predicted target state and predicted robot state given the control $\mathbf{a}_i(t)$.

The first cost function, $J_1(t)$, penalizes the deviation of the target position from the FoV center [18]:

$$J_1(t) = \exp\left(\sqrt{\rho_{in}^{d}(t+1)^2 + \rho_{in}^{\delta}(t+1)^2}\right), \qquad (1)$$

$J_1(t) < \exp(1)$ ensures that the target locates within the FoV. The minimum, $J_1(t) = 1$, is achieved when the target is at the center of the FoV.

The second cost function, $J_2(t)$, penalizes accelerations/decelerations induced by $J_1(t)$:

$$J_2(t) = \exp\left(\frac{\|\mathbf{v}_i(t+1) - \mathbf{v}_i(t)\|}{v_{\max} + \|\mathbf{v}_i(t)\|}\right), \qquad (2)$$

where $v_{\max} + \|\mathbf{v}_i(t)\|$ is the maximum speed difference of a robot between two consecutive time steps. The minimum, $J_2(t) = 1$, is achieved when the robot does not change velocity.

The control is computed by minimizing with brute force the two cost functions:

$$\mathbf{a}_i(t) = \text{argmin}_{\mathbf{a}}\left(\lambda J_1(t) + (1-\lambda) J_2(t)\right), \qquad (3)$$

where $\lambda \in [0.5\ ,\ 1]$. λ starts at 0.5 as the main objective is to maintain the target centered in the FoV. We vary λ based on whether the robot has its target inside the FoV. When the robot loses its target from the FoV, the objective is to recapture it as soon as possible, and we therefore set $\lambda = 1$ without accounting for motion smoothness. When the robot has its target inside the FoV, the objective is to maintain the target centered in the FoV with smooth motion. Based on a sensitivity test on λ, we experimentally set $\lambda = 0.6$ as it provides the best trade-off between the decrement of $J_1(t)$ and the increment of $J_2(t)$.

3.3. Accessible collision-avoiding velocity set

Let us consider[1] a pair of robots c_i and c_j at position $\mathbf{p}_i$ and $\mathbf{p}_j$ respectively, aiming to achieve their preferred velocity $\mathbf{v}_i^*$ and $\mathbf{v}_j^*$ (Fig. 2(a)). Let $\mathbf{O}_{i,j}^{\tau}(\mathbf{0})$ be the velocity obstacle of c_i induced by c_j, which is the set of relative velocities of c_i with respect to c_j that can lead to a collision within a time horizon τ (Fig. 2(b)):

$$\mathbf{O}_{i,j}^{\tau}(\mathbf{0}) = \{\mathbf{v} | \exists t \in [0, \tau], \|t\mathbf{v}\| \geqslant \left\|\mathbf{p}_{ij} - 2r\right\|\}, \qquad (4)$$

where $\mathbf{p}_{ij} = \mathbf{p}_j - \mathbf{p}_i$ is the relative position of c_j with respect to c_i. The set of collision-avoiding relative velocities for c_i to avoid c_j in τ time horizon, $\mathbf{A}_{i,j}^{\tau}(\mathbf{0})$, is therefore:

$$\mathbf{A}_{i,j}^{\tau}(\mathbf{0}) = \{\mathbf{v} | \mathbf{v} \notin \mathbf{O}_{i,j}^{\tau}(\mathbf{0})\}. \qquad (5)$$

[1]To simplify the notation we omit t from here.

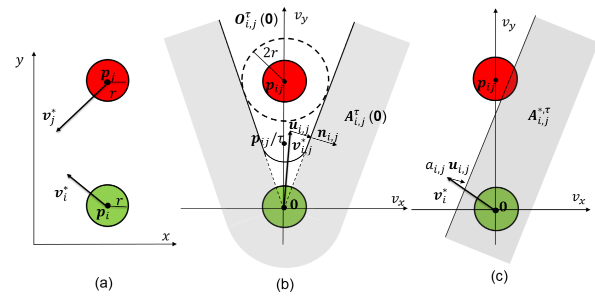

Figure 2. Optimal reciprocal collision-avoiding velocities. (a) Robot c_i and c_j with radius r at $\mathbf{p}_i$ and $\mathbf{p}_j$ with preferred velocity $\mathbf{v}_i^*$ and $\mathbf{v}_j^*$, respectively. (b) The gray shadow area indicates $\mathbf{A}_{i,j}^{\tau}(\mathbf{0})$, the relative velocities of c_i that are collision-avoiding to c_j in τ time steps. $\mathbf{u}_{i,j}$ is the minimal velocity for $\mathbf{v}_{i,j}^*$ to get out of the velocity obstacle and $\mathbf{n}_{i,j}$ is the outward normal at $\mathbf{v}_{i,j}^* + \mathbf{u}_{i,j}$. (c) The gray shadow area indicates the velocities of c_i that are optimal reciprocal collision-avoiding to c_j in τ time steps ($\mathbf{A}_{i,j}^{*,\tau}$ of c_i) when c_i shares $a_{i,j}$ responsibility to avoid c_j.

Reciprocal collision avoidance is possible when c_i and c_j choose to move at $\mathbf{v}_i$ and $\mathbf{v}_j$ respectively and $\mathbf{v}_i \in \mathbf{A}_{i,j}^{\tau}(\mathbf{v}_j)$ and $\mathbf{v}_j \in \mathbf{A}_{j,i}^{\tau}(\mathbf{v}_i)$ [17].

ORCA [17] defines the set of velocities that is not only reciprocal collision-avoiding but also close to the preferred velocity. Let $\mathbf{v}_{i,j}^* = \mathbf{v}_i^* - \mathbf{v}_j^*$ be the relative preferred velocity of c_i with respect to c_j. Let $\mathbf{u}_{i,j}$ be the vector starting from $\mathbf{v}_{i,j}^*$ to the closest point on the boundary of $\mathbf{O}_{i,j}^{\tau}(\mathbf{0})$ (Fig. 2(b)). $\mathbf{u}_{i,j}$ is the minimal relative velocity change between c_i and c_j to avoid collisions within τ. $\mathbf{n}_{i,j}$ is the outward plane normal at $\mathbf{v}_{i,j}^* + \mathbf{u}_{i,j}$.

Each robot shares a partial responsibility for collision avoidance. Let $a_{i,j}$ and $a_{j,i}$ be the responsibility c_i and c_j take to avoid each other and $a_{i,j} + a_{j,i} = 1$. The responsibility $a_{i,j}$ indicates how much c_i will compensate $\mathbf{u}_{i,j}$ in order to shift $\mathbf{v}_{i,j}^*$ out of the velocity obstacle. The set of optimal reciprocal collision-avoiding velocity for c_i to avoid c_j in τ time steps is:

$$\mathbf{A}_{i,j}^{*,\tau} = \{\mathbf{v} | \mathbf{v} - (\mathbf{v}_i^* + a_{i,j}\mathbf{u}_{i,j}) \cdot \mathbf{n}_{i,j} \leqslant 0\}. \qquad (6)$$

$\mathbf{A}_{i,j}^{*,\tau}$ are the velocities that lie at the half-plane (the plane in gray shadow in Fig. 2(c)) in the direction of $\mathbf{n}_{i,j}$, after shifting $\mathbf{O}_{i,j}^{\tau}(\mathbf{0})$ by $\mathbf{v}_i^* + a_{i,j}\mathbf{u}_{i,j}$. In the same way, we can construct $\mathbf{A}_{j,i}^{*,\tau}$ for c_j.

Let $\mathbf{V}_i$ be the set of velocities that are accessible under the speed/acceleration limits and nonholonomic constraint. Each robot c_i estimates the pair-wise $\mathbf{A}_{i,j}^{*,\tau}, \forall c_j \in C_i^A$, where C_i^A refers to the set of robots within the avoidance range of c_i. The final set of accessible collision-avoiding

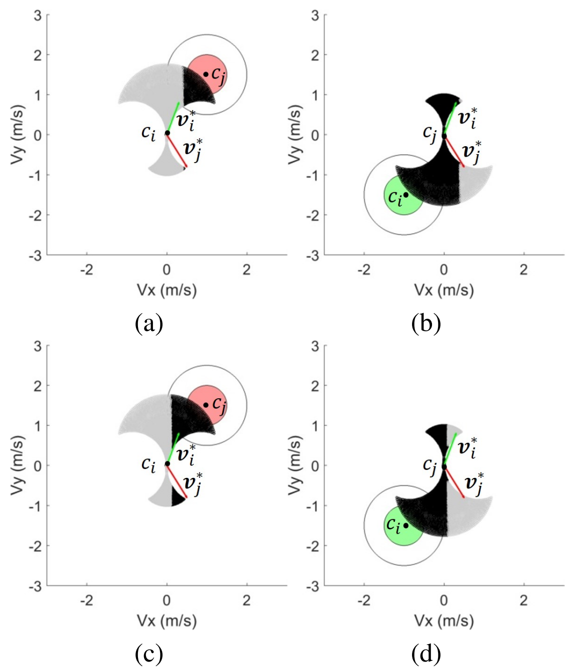

Figure 3. The accessible collision-avoiding velocity set (gray area) is influenced by the pair-wise responsibility sharing. $\mathbf{v}_i^*$ and $\mathbf{v}_j^*$ are the preferred velocities of c_i and c_j respectively. The collision-avoiding velocities of (a) c_i induced by c_j and (b) c_j induced by c_i with equal responsibility, where c_j is likely to experience empty-set if surrounded by other robots. The collision-avoiding velocities of (c) c_i induced by c_j and (d) c_j induced by c_i with adaptive responsibility.

velocities for c_i, $\mathbf{A}_i^{*,\tau}$, is

$$\mathbf{A}_i^{*,\tau} = \left(\bigcap_{c_j \in C_i^A} \mathbf{A}_{i,j}^{*,\tau} \right) \cap \mathbf{V}_i. \tag{7}$$

When neighboring robots leave no space for a robot to move (i.e. $\mathbf{A}_i^{*,\tau} = \emptyset$), the robot may lose its target. A robot should therefore reduce the occurrences of empty-set cases.

3.4. Adaptive responsibility sharing

When a robot has a small set of accessible collision-avoiding velocities induced by a neighboring robot, it is more likely to incur an empty-set case if it is surrounded by additional robots. The choice of the *responsibility* between a pair of robots influences the total number of accessible collision-avoiding velocities (see Fig. 3).

Let us re-write the notation $\mathbf{A}_{i,j}^{*,\tau}$ to $\mathbf{A}_{i,j}^{*,\tau}(a)$ to indicate the dependence on the responsibility. Our objective is to maximize the *total* number of optimal reciprocal collision-avoiding velocities, while maintaining the *fairness*, i.e. a balanced distribution of accessible collision-avoiding velocities between a pair of robots.

We compute the percentage of optimal reciprocal collision-avoiding velocities for c_i induced by c_j as $\rho_{i,j}^{*,\tau}(a) = \frac{|\mathbf{A}_{i,j}^{*,\tau}(a)|}{|\mathbf{V}_i|}$, where $|\cdot|$ is the cardinality of a set.

The responsibility for c_i to avoid c_j is computed by maximizing the average percentage of optimal reciprocal collision-avoiding velocities while accounting for fairness:

$$a_{i,j} = \operatorname{argmax}_a \left(f_{i,j}(a) \frac{\rho_{i,j}^{*,\tau}(a) + \rho_{j,i}^{*,\tau}(1-a)}{2} \right), \tag{8}$$

where $f_{i,j}(a)$ indicates the fairness of the percentage of optimal reciprocal collision-avoiding velocities between a pair of robots. In order to have a continuous and bounded fairness measure to combine with the average ratio, we adopt the Jain's fairness measure [9]:

$$f_{i,j}(a) = \frac{\left(\rho_{i,j}^{*,\tau}(a) + \rho_{j,i}^{*,\tau}(a)\right)^2}{2(\rho_{i,j}^{*,\tau}(a)^2 + \rho_{j,i}^{*,\tau}(a)^2)}. \tag{9}$$

Note that the responsibility can be a negative value as long as the responsibilities shared by a robot pair sums to 1.

We compute the optimal $a_{i,j}$ by searching in a discretized space bounded by $[-a_{\max},\ a_{\max}+1]$. We compute the adaptive responsibility only if a robot has less than half of accessible collision-avoiding velocities when sharing equal responsibility. The adaptive responsibility $a_{j,i}$ for c_j to avoid c_i is computed in the same way. We can guarantee $a_{i,j} + a_{j,i} = 1$ due to the symmetric property of the objective function.

Note that this scheme reduces the chances of the occurrence of empty-set cases but does not prevent empty-set cases. When an empty-set case occurs, we use a simple braking-recompute strategy: the robot with an empty set stops and all other robots recompute the collision-avoiding velocity sets. The process is repeated until no robot is in an empty-set situation.

4. Validation

We validate the proposed framework, A-ORCA, in terms of the influence of collision avoidance on target viewing performance and compare it with the avoidance strategies based on RVO [16] and ORCA [17].

We quantify the empty-set occurrences, ε, as the empty-set ratio

$$\varepsilon = \frac{1}{MT} \sum_{i=1}^{M} \sum_{t=1}^{T} b_i^e(t), \tag{10}$$

where T is the experiment time steps, M is the number of robots and $b_i^e(t)$ is a binary value indicating whether the set of accessible collision-avoiding velocities is empty ($A_i^{*,\tau}(t) = \emptyset$). Similarly, we quantify the target

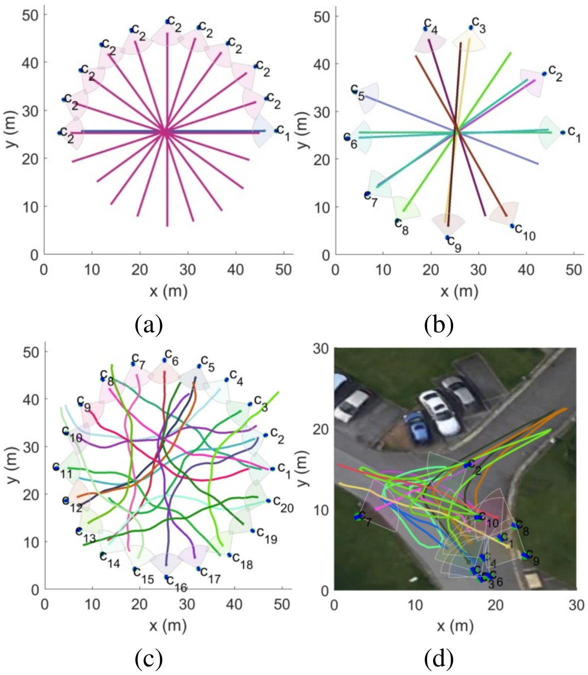

Figure 4. The four scenarios used for validation. The color of the trajectory indicates the target ID. (a) Scenario I with a pair of targets intersecting at different angles ranging from 0.1π to π with step size 0.1π. The target tracked by c_1 is fixed at angle 0 and we vary the angle of the target tracked by c_2. (b) Sample Scenario II with 10 targets with randomly distributed intersecting angles between a pair of neighbors. (c) Sample Scenario III with 20 targets that intersect at random times and locations. (d) Scenario IV with 10 trajectories extracted from a video sequence of the PET2009 dataset.

viewing performance of robots, η, by the viewing ratio

$$\eta = \frac{1}{MT} \sum_{i=1}^{M} \sum_{t=1}^{T} b_i^v(t), \tag{11}$$

where $b_i^v(t)$ is a binary value for the presence of target within FoV. Results are averaged over 10 independent runs.

We use a real scenario and design three synthetic scenarios to investigate the influence of the *intersecting angles* between targets and the *number* of intersecting targets on the methods (Fig. 4). Targets are initialized along a circle with their initial velocity heading towards the center of the circle. Each target moves at a speed of $1m/s$ with additional Gaussian noise bounded within $[-0.1m/s, 0.1m/s]$. Scenario I increases the intersecting angle between a pair of targets from 0.1π to π with step size 0.1π (Fig. 4(a)). Scenario II increases the number of targets from 2 to 10 and randomly sets the intersecting angle between two neighboring targets (Fig. 4(b)). Scenario III increases the number of targets from 2 to 30 and creates sparser target intersections by varying the velocity direction over time (Fig. 4(c)).

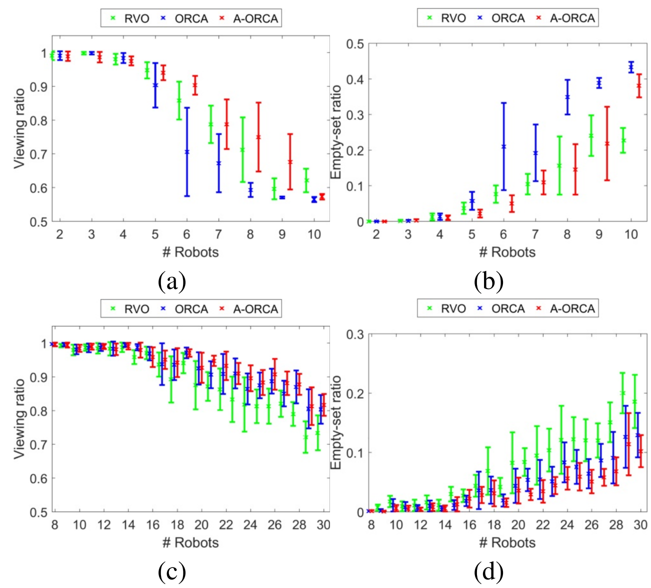

Figure 5. Target viewing ratio, η, and empty-set ratio, ε, for different collision avoidance methods when increasing the number of robots. Top row: viewing ratio (a) and empty-set ratio (b) under Scenario II. Bottom row: viewing ratio (c) and empty-set ratio (d) under Scenario III with eight or more robots (all methods achieve full viewing ratio without empty-set with less than eight robots).

Additionally, we test the methods in Scenario IV with real people trajectories extracted from the PETS2009 dataset[2]. The trajectories of 10 people of $60s$ duration is extracted from the S2L1 sequence and map to the ground plane with the provided camera calibration (Fig. 4(d)).

In all scenarios, robots follow a car-like model and are initialized with their target centered in their camera FoV. Each camera has the viewing angle $\phi = 90°$ and viewing range $r_v = 5m$. We set the robot avoidance range to $2v_{\max}$ as it is the worst case for a collision between a pair of robots in the next time step. We set $\tau = 6$ for Scenario II and $\tau = 3$ for the other scenarios. A larger τ allows for earlier aversion when robots are densely intersecting (Scenario II), and improves the viewing ratio.

Under Scenario I, all methods successfully avoid collisions and robots keep their target in their FoV regardless of the variation in target intersecting angles. Under Scenario II, the viewing ratio deteriorates as the number of robots, M, increases. When $M > 7$, robots using ORCA achieve a limited average viewing ratio (60%) with a small standard deviation because all robots get stuck at the intersecting position (Fig. 5(a)). Robots using A-ORCA achieves a higher viewing ratio than those using ORCA, as empty-set occurrences are reduced by almost half (Fig. 5(b)). The pair-wise optimization does not provide guarantee for non-empty sets, and this explains that A-ORCA fails with $M = 10$. In-

[2] http://www.cvg.reading.ac.uk/PETS2009. Last accessed: 07/04/2017

terestingly, robots using RVO may achieve higher viewing ratios than ORCA. Compared to ORCA, RVO has more restrictive velocity constraints which prevent robots from getting too close to each other and therefore provides more space for robots to avoid each other and recapture their target (see [1]).

Under Scenario III, ORCA-based methods outperform RVO as the target trajectories require less effort for avoidance (Fig. 5(c)). Although A-ORCA reduces the empty-set ratio in this scenario, it may achieve a slightly worse average viewing ratio compared to ORCA when $13 \leqslant M \leqslant 16$. This is because A-ORCA adapts the responsibility sharing among robots when there is an empty-set case, in order to achieve a non-empty set of collision-avoiding velocities. A-ORCA may end up with a new velocity that causes the robot to temporally lose its target from the FoV, while ORCA may simply stop and move at next time step without losing the target.

Finally, in Scenario IV people cross each other, walk in parallel and stay together without moving (see [1]). No collision occurred and the viewing ratios (empty-set ratios) achieved by RVO, ORCA and A-ORCA are 0.96 (0.032), 0.98 (0.018) and 0.96 (0.017), respectively. The results are consistent with those obtained with the synthetic trajectories when $M = 10$ (Fig. 5(c)).

5. Conclusion

We presented an ORCA-based framework for active visual tracking in multi-robot scenarios. We proposed a controller that accounts for both viewing performance and motion smoothness. We improved ORCA with an adaptive scheme to optimally set the pair-wise responsibility to reduce the empty-set cases under target intersecting scenarios. This solution leads to longer target viewing.

As future work, we will include a target avoidance functionality and validate the framework with real robots.

References

[1] http://www.eecs.qmul.ac.uk/~andrea/aorca.html.

[2] J. Alonso-Mora, A. Breitenmoser, P. Beardsley, and R. Siegwart. Reciprocal collision avoidance for multiple car-like robots. In *Proc. of IEEE Int'l Conf. on Robotics and Automation*, pages 360–366, Saint Paul, MN, May 2012.

[3] J. Alonso-Mora, T. Naegeli, R. Siegwart, and P. Beardsley. Collision avoidance for aerial vehicles in multi-agent scenarios. *Autonomous Robots*, 39:101–121, Jan 2015.

[4] D. Bareiss and J. van den Berg. Generalized reciprocal collision avoidance. *Int'l J. of Robotics Research*, 34(12):1501–1514, Oct 2015.

[5] A. Best, S. Narang, and D. Manocha. Real-time reciprocal collision avoidance with elliptical agents. In *Proc. of IEEE Int'l Conf. on Robotics and Automation*, pages 298–305, Stockholm, Sweden, May 2016.

[6] S. Curtis, B. Zafar, A. Gutub, and D. Manocha. Right of way: Asymmetric agent interactions in crowds. *Visual Computer*, 29(12):1277–1292, 2013.

[7] P. Fiorini and Z. Shiller. Motion planning in dynamic environments using velocity obstacles. *Int'l Journal of Robotics Research*, 17:760–772, 1998.

[8] A. E. Guevara, A. Hoak, J. T. Bernal, and H. Medeiros. Vision-based self-contained target following robot using bayesian data fusion. In *Proc. of Int'l Symp. on Visual Computing*, pages 1–12, Las Vegas, USA, Dec 2016.

[9] R. Jain, D. Chiu, and W. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. *Computing Research Repository*, cs.NI/9809099, 1998.

[10] M. Luber, L. Spinello, J. Silva, and K. O. Arras. Socially-aware robot navigation: A learning approach. In *Proc. of IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, pages 902–907, Vilamoura, Algarve, Portugal, Oct 2012.

[11] J. Ondřej, J. Pettré, A.-H. Olivier, and S. Donikian. A synthetic-vision based steering approach for crowd simulation. *ACM Trans. Graph.*, 29(4):123:1–123:9, July 2010.

[12] J. Pestana, J. Sanchez-Lopez, S. Saripalli, and P. Campoy. Computer vision based general object following for gps-denied multirotor unmanned vehicles. In *Proc. of American Control Conference*, pages 1886–1891, Portland, OR, June 2014.

[13] F. Poiesi and A. Cavallaro. Distributed vision-based flying cameras to film a moving target. In *Proc. of IEEE Int'l Conf. on Intelligent Robots and Systems*, pages 2453–2459, Hamburg, Germany, Sep 2015.

[14] J. Snape, J. van den Berg, S. J. Guy, and D. Manocha. Independent navigation of multiple mobile robots with hybrid reciprocal velocity obstacles. In *Proc. of IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, pages 5917–5922, St. Louis, USA, Oct 2009.

[15] J. Snape, J. van den Berg, S. J. Guy, and D. Manocha. Smooth and collision-free navigation for multiple robots under differential-drive constraints. In *Proc. of IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, pages 4584–4589, Taipei, Oct 2010.

[16] J. van den Berg, M. Lin, and D. Manocha. Reciprocal velocity obstacles for real-time multi-agent navigation. In *Proc. of IEEE Int'l Conf. on Robotics and Automation*, pages 1928–1935, Pasadena, CA, May 2008.

[17] J. van den Berg, G. Stephen J., L. Ming, and M. Dinesh. Reciprocal n-body collision avoidance. In *Proc. of Int'l Symp. on Robotics Research*, volume 70, pages 3–19. Springer Berlin Heidelberg, 2009.

[18] Y. Wang and A. Cavallaro. Prioritized target tracking with active collaborative cameras. In *Proc. of Int'l Conf. on Advanced Video and Signal-Based Surveillance*, pages 131–137, Colorado Springs, Aug 2016.

[19] H. Wei, W. Lu, P. Zhu, G. Huang, J. Leonard, and S. Ferrari. Optimized visibility motion planning for target tracking and localization. In *Proc. of IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, pages 76–82, Chicago, IL, Sept 2014.