

Following a Large Unpredictable Group of Targets Among Obstacles

Christopher Vo and Jyh-Ming Lien

Department of Computer Science, George Mason University,
4400 University Drive MSN 4A5, Fairfax, VA 22030 USA
{cvo1, jmlien}@gmu.edu
<http://masc.cs.gmu.edu>

Abstract. Camera control is essential in both virtual and real-world environments. Our work focuses on an instance of camera control called target following, and offers an algorithm, based on the ideas of *monotonic tracking regions* and *ghost targets*, for following a large coherent group of targets with unknown trajectories, among known obstacles. In multiple-target following, the camera’s primary objective is to follow and maximize visibility of multiple moving targets. For example, in video games, a third-person view camera may be controlled to follow a group of characters through complicated virtual environments. In robotics, a camera attached to robotic manipulators could also be controlled to observe live performers in a concert, monitor assembly of a mechanical system, or maintain task visibility during teleoperated surgical procedures. To the best of our knowledge, this work is the first attempting to address this particular instance of camera control.

Keywords: motion planning, camera planning, target following, group motion monitoring, monotonic tracking regions, ghost targets

1 Introduction

In *multiple-target following*, the camera’s primary objective is to follow and maximize visibility of multiple moving targets. Multiple-target following is essential in both virtual and real-world environments. For example, in video games, a third-person view camera may be controlled to follow a group of characters through complicated virtual environments. In robotics, a camera attached to robotic manipulators could be controlled to observe a swarm of mobile robots, or live performers in a concert; monitor assembly of a mechanical system; or maintain task visibility during teleoperated surgical procedures.

In general, it is difficult for a user to manually control the camera while also concentrating on other critical tasks. Therefore, it is desirable to have an *autonomous camera system* that handles the camera movement. This paper focuses on an instance of camera control called *multiple-target following*, and offers an algorithm for *autonomous* following a large coherent group of 10 ~ 100 targets (such as a crowd or a flock) with unknown trajectories, among known obstacles.

The camera control problem has been studied extensively in both robotics and computer graphics because of its broad applications, such as dynamic data visualization [21], robotic and unmanned vehicle teleoperation [13], and video games [16]. Unfortunately, many of these methods are not applicable to follow a large group of targets in real-time *among obstacles*. For example, there is a large body of work in robotics, where researchers have studied similar problems such as pursuit and evasion, visual servoing [6] and cooperative multi-robot observation of multiple moving targets (CMOMMT) [19]. However, these strategies usually apply to environments with sparse or no obstacles. While there exist methods that do consider occlusion, e.g. [15], they still only consider situations where the target’s trajectory is known *a priori*, or are only applicable to follow a few (2 or 3) targets [14,17]. Many researchers have also considered the case where both trajectory and environment are unknown [11,1]. The main idea of these motion planners is to greedily minimize the *escaping risk* or maximize the *shortest escaping distance* of the target. In all of these methods, the camera trajectory can be computed efficiently, but is usually sub-optimal because only local information about the environment is considered, and the time horizon for planning is very short. Moreover, to the best of our knowledge, these motion strategies are all designed to track a single target.

In computer graphics, much work on camera planning is script-based [5], purely reactive [12,6], or mostly focuses on problems with predefined target trajectories [2,9]. Many of these methods are based on constraint solving, objective satisfaction or both. They are mostly designed for offline use and take a long time (usually seconds) to find a single camera placement.

The main goal of this paper is to provide an initial investigation into this important problem. Due to the nature of the aforementioned applications (video games, mobile robot swarms, virtual prototyping and group control), our investigation will focus on following a coherent group of targets by a single camera. Maintaining the visibility of a coherent group in some aspects is easier than tracking a single target because there is more than one target that the camera can follow. However, tracking a group is more difficult if the camera needs to maximize the number of visible targets over time. Difficulty also stems from the fact that a group can assume different shapes (e.g., forming a long line in a narrow corridor and a blob in an open area), clutter around the obstacles, or even split into multiple sub-groups (for a short period of time).

Our Work and Main Contributions. In this paper, we present a motion strategy that allows a single camera to robustly follow, at interactive rates, a large group (e.g., with 100 targets) among obstacles. In the aforementioned applications, some information about the environment is usually available. We believe that, given this information about the environment, the planner should be able to perform deeper *lookahead* in the search space and therefore provide better real-time camera following strategies even when the motion of the targets is unknown. The main idea is to preprocess the given environment offline to generate a data structure called *monotonic tracking regions* (MTRs) (defined later in Section 4.1) that can be used to assist real-time planning. This representation

allows the camera to plan more efficiently by reducing the possible target movements to a smaller, discrete space. This significant increase in efficiency allows us to generate and evaluate multiple alternative plans in real-time. Our method also uses target coherency to better predict target positions.

In addition to the MTR-based motion strategy, we also present three new strategies (in Section 5) extended from the existing single-target techniques. We present reactive, sampling-based [3], and escaping-risk-based [18] methods. Our experimental results show that MTR-based method performs significantly better than the other strategies, especially in the environments with small obstacles and sinuous tunnels. A preliminary version of this work on both searching and tracking large group can be found at [22]. To the best of our knowledge, this work is the first attempting to address this particular instance of camera control.

2 Related Work

There exist many methods for following a target with known trajectory, such as work done by LaValle et al. [15] using dynamic programming, and by Goemans and Overmars [10] using probabilistic motion planner. In this section, we will review strategies for tracking a target with unknown trajectory. There are also extensive studies on multiple object tracking, e.g., [20], in which the goal is to maintain a belief of where the target(s) are using Kalman or particle filters. On the contrary, our goal is to maintain the visibility of targets. From our review, no method has focused on the problem of following a large group.

There exists some work considering tracking targets with unknown trajectories in a known environment [17,4]. The general idea is to partition the space into non-critical regions in which the camera can follow the target without complex *compliant motion*, that is, rotating the line of visibility between the camera and the target around a vertex of an obstacle. The main benefit of this line of work is the ability to determine the *decidability* of the camera tracking problem [4].

Recently, Li and Cheng [16] have proposed a real-time planner that tracks a target with unknown trajectory. Their main ideas include a *budgeted roadmap method* with lazy evaluation. Geraerts [8] proposed the idea of using the *corridor map* to track a single target with the corridors.

In computer graphics, camera planning is often viewed as a constraint satisfaction problem, and so there have been several attempts to represent the problem so that it can be solved efficiently with constraint satisfaction techniques. For example, several works use the idea of *screen space* or *image space* constraints, e.g., Gleicher and Witkin [9]. There are a number of works which involve the use of metaheuristics to compute optimal positions or trajectories for the camera. For example, Drucker and Zeltzer [7] used an A* planner to compute a camera path. Along the path, the orientation of the camera is then solved frame-by-frame to satisfy given constraints.

Some studies have focused on developing reactive behaviors for real-time camera motion. For example, Halper et al. [12] introduced a camera planner that predicts state based on the past trajectory and acceleration. They also

proposed the idea of PVR (potential visibility region) for visibility computation, and a pipelined constraint solver.

3 Preliminaries

In this section, we formally define the problem that we attempt to solve and notations used throughout the paper. We assume that the workspace is populated with known obstacles represented by polygons. These polygons are the projection of 3D objects that can potentially block the camera’s view. This projection essential reduces our problem to a 2D workspace. We further assume that, initially, at least one of the targets is visible by the camera C , and, during the entire simulation, the targets T exhibit certain degree of coherence in their motion, and it is also possible that T can split into multiple subgroups for a short period of time (similar to a flock of birds). The targets are either controlled by the user or by another program, so the trajectories of the targets are not known in advance. However we assume that the size of T and T ’s maximum (linear) velocity v_T^{max} are known. The position $x_\tau(t)$ and the velocity $v_\tau(t)$ of a target $\tau \in T$ at time t are known only if τ is visible by the camera.

The camera C also has a bounded linear velocity v_C^{max} . The exact configuration of this view range at time t , denoted as $\mathcal{V}_C(t)$, is defined by the camera’s view direction $\theta_C(t)$ and location $x_C(t)$. The position x_C of the camera is simply governed by the following equation: $x_C(t + \Delta t) = x_C(t) + \Delta t \cdot v_C(t)$, where $v_C(t)$ is the camera’s velocity at time t .

Given the positions of the targets and the position of the camera, one can compute the camera’s view direction so that the number of targets inside the view range is maximized. Therefore, the problem of target following then is reduced to find a sequence of velocities $v_C(t)$:

$$\arg \max_{v_C(t)} \left(\sum_t \text{card}(\{T' \subset T \mid X_{T'}(t) \subset \mathcal{V}_C(t)\}) \right), \quad (1)$$

subject to the constraints that, $v_C(t) \leq v_C^{max}, \forall t$, and $x_C(t)$ is collision free.

The main ideas of the proposed method are to (1) identify regions with simple (monotonic) topological feature so that the planner can repetitively use the same data structure and strategy to follow the target group, and (2) utilize the fact the targets form a coherent group.

4 Monotonic Tracking Regions

The first step of the proposed method decomposes the environment into a set of *monotonic tracking regions* (MTRs). These regions usually look like tunnels and may overlap with each other. Intuitively, in these tunnel-like regions, the camera can *monotonically* maintain the visibility by moving forward or backward along a trajectory that *supports* the tunnel. More specifically, the main property

MTR is that each MTR is topologically a *linear subdivision* so that the problem of camera following in an MTR can be represented as a *linear programming problem*. Note that such an MTR needs not to be convex or star shaped, and, in fact, it can have an arbitrary number of turns (like a sinuous tunnel). Moreover, MTR decomposition usually creates much fewer components than convex or star-shaped decompositions but, as we will see later, still provide similar functionality in target tracking. More precise definition of MTR and the process of computing these regions will be given in Section 4.1. In Sections 4.2 and 4.3, we will discuss how to track the target in MTR and ghost targets, respectively.

4.1 Build Monotonic Tracking Regions

Definition. We let a region M_π be a 2D generalized cylinder defined with respect to a *supporting path* π . We say π is a supporting path of M_π if every point $x \in M_\pi$ can see a subset of π . Because of this property, M_π can essentially be viewed as a linear object and the camera can see every point in M_π by moving along π . Specifically, we define M_π as:

$$M_\pi = \{x \mid \exists y \in \pi \text{ s.t. } \overline{xy} \subset C_{free}\}, \quad (2)$$

where \overline{xy} is an open line segment between x and y , and C_{free} is the free space (i.e., the area without obstacles). Furthermore, we define the subset of π visible by x as: $V_\pi(x) = \{y \in \pi \mid \overline{xy} \subset C_{free}\}$. Note that $V_\pi(x)$ can have one or multiple connected components. Finally, we say a region $M_\pi \in C_{free}$ is an MTR supported by π if

$$|V_\pi(x)| = 1, \forall x \in M_\pi, \quad (3)$$

where $|X|$ is the number of connect components in a set X . Because each $x \in M_\pi$ can see only an interval of π , we can compactly represent the visible region (called visibility interval) of x as a tuple $V_\pi(x) = (s, t), 0 \leq s \leq t \leq 1$, if we parameterize π from 0 to 1. Fig. 1 shows an example of MTR and its supporting path π .

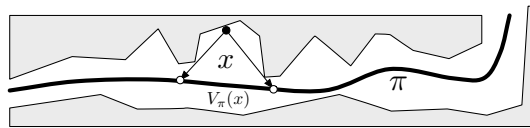


Fig. 1. An example of monotonic tracking regions defined by π

With the definition in hand, our task here is to first find a set of supporting paths whose MTRs that will cover C_{free} , and, next, from a given path π , we compute the associated MTR and the visibility interval $V_\pi(x)$ for every point x in the MTR. We will describe these two steps in detail next.

Constructing supporting paths. Our strategy here is to find the homotopy groups G of the C_{free} . We propose to use the medial axis (MA) of the C_{free}

to capture G because of its several interesting relationships with MTRs. First of all, we can show that the *retraction region* of every edge π on the MA forms an MTR (supported by π).

Lemma 1. *The retraction region $R \subset C_{free}$ of an MA edge forms an MTR.*

Proof. Let π be an edge on the medial axis MA of C_{free} . The retraction region $R \subset C_{free}$ of π is simply a set of points that can be continuously retracted to π by a retraction function $r : R \rightarrow \pi$ [23]. By definition, given an arbitrary point $x \in R$, the largest circle c centered at the point $r(x)$ with x on c 's boundary must be empty. Therefore, it follows naturally that each point in R must be able to see at least a point on π .

Now we briefly show that each point x can only see a consecutive region of π . We prove this by contradiction. Assuming that x can see multiple intervals of π . This means that there must be an obstacle between x and π . However, this contradicts the fact that x can be retracted to π in a straight line. \square

Therefore, the supporting paths are simply constructed by extracting the edges from the MA of a given environment.

Constructing mtrs. Given an edge π of MA, its retraction region R forms an MTR supported by π . However, simply using R as π 's MTR can be overly conservative. The set of points that satisfy Eq. 2 and 3 is usually larger than R . To address this issue, we iteratively expand R by considering the points adjacent to R until no points can be included without violating the definition of MTR.

Next, we compute the visibility interval for every point in an MTR. The brute force approach that computes the visibility interval for each point from scratch is no doubt time consuming. To speed up the computation, our approach is based on the following observation.

Observation 41 *If x and x' are (topological) neighbors, and x is further away from π than x' is, then $V_\pi(x) \subset V_\pi(x')$.*

For example, in Fig. 1, imagining a point x' below x , x' can see a larger interval of π than x does. That is if we can compute the visibility intervals $V_\pi(x')$ for all the points x' on π , then we should be able to obtain the visibility intervals for x that are Δd away from x' by searching inside $V_\pi(x')$.

Dominated mtrs. The exact MA of a given environment can contain small (and in many cases unnecessary) features and result in many small MTRs. In many cases, these MTRs are unnecessary and should be removed. This is when an MTR is dominated by another MTR. We say MTR M' is dominated by another MTR M if $M' \subset M$. In our implementation, we use an approximate MA [23] to avoid small features, and then identify and remove dominated MTRs.

4.2 Follow the Targets in an MTR

The motivating idea behind decomposing the environment into a set of MTRs is that the target following problem in MTR can be solved much easily than that in the original environment. In fact, as we will see in this section, the camera can solve a long time horizon plan in MTR using linear programming.

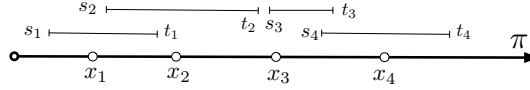


Fig. 2. Make predictions for the next $h = 4$ future steps.

Follow a single target. To simplify our discussion, we will first describe how a single target can be tracked in MTR. Let $x_\tau(t)$ be the current position of the target τ . Since we know the current speed of the target, we can estimate the positions $x_\tau(t + \Delta t)$ in the next time step. In order to keep the target in the view, the camera's next position $x_C(t + \Delta t)$ must be: $x_C(t + \Delta t) \in V_\pi(x_\tau(t + \Delta t))$. Note that this estimation can be applied to an arbitrary value of Δt . However, when Δt is bigger, the position of the target becomes less accurate.

Let $I_i = V_\pi(x_\tau(t + i \cdot \Delta t)) = (s_i, t_i)$. Here i is an integer from 1 to h , where h is the user-defined time horizon. Recall that both s_i and t_i are parameters on the parameterized path π . In order to follow the target for h steps, the planner needs find a sequence of camera locations x_i from a sequence of parameterized intervals such that every point x_i is in its corresponding interval I_i without violating the constraint on the camera's max speed (see Fig. 2). This can be done by solving a h dimensional linear programming problem:

$$\begin{aligned} \min \quad & t_h - x_h \\ \text{s.t.} \quad & s_i \leq x_i \leq t_i \\ & 0 \leq (x_{i+1} - x_i) \leq \frac{v_C^{max}}{|\pi|}, \forall x_i, \end{aligned} \quad (4)$$

where $v_C^{max}/|\pi|$ is the maximum *normalized* distance that the camera can travel on π . Finally, the camera's future locations are simply $x_C(t + \Delta t \cdot i) = \pi(x_i)$.

Note that the rationale behind the minimization of $(t_h - x_h)$ is that when the target moves further away beyond h steps in the future, the camera will have better chance of keeping the target in the view when it is located closer to t_h along the path π . We call the above linear programming problem the *canonical following problem*. Solving a canonical following problem can be done efficiently since h is usually not large ($h = 20$ is used in our experiments) given that modern linear programming solvers can handle thousands of variables efficiently.

It is possible that the linear programming problem has no feasible solution. We reduce the plan horizon iteratively until a solution is found.

Follow multiple targets. Now, we extend this canonical following problem to handle multiple targets T . Let $x_T(t)$ be the current positions of T . Similar to the case of a single target, we estimate the positions $x_T(t + \Delta t)$ in the next time step. In order to see a least one target, the camera must move so that

$$x_C(t + \Delta t) \in I(\Delta t) = \bigcup_{x \in x_T(t + \Delta t)} V_\pi(x) .$$

To simplify our notation, let $I_i = I_i(i \cdot \Delta t) = (s_i, t_i)$. By placing the camera in I_i , we can guarantee that at least one target is visible. However, our goal is to

maximize the number of visible targets, at least over the planning horizon. To do so, we segment I_i into j sub-intervals I_i^j , each of which can see n_i^j targets. Then our goal is to pick a sub-interval from each I_i so that the total number of visible targets is maximized while still maintaining the constraint that the minimum distance between I_i^j and I_{i+1}^k is smaller than $v_T^{max} \Delta t$. Fortunately, this optimization problem can be solved greedily by iteratively adding the sub-interval with the largest n_i^j without violating the constraint. Once the sub-interval from each I_i is identified, the problem of finding the camera positions is formulated as a linear programming problem in the same way as Eq. 4.

When the targets and the camera moves (but still in the same MTR), we may need to repetitively solve the canonical following problem. Instead, we use *lazy update*. That is, we only update the old solution if it cannot satisfy all the constraints in Eq. 4.

4.3 Ghost targets

Our planner takes advantage of the fact that the camera is following a group of somewhat coherent targets. When the number of the targets visible by the camera is smaller than the total number of the targets, the planner will generate a set of *ghost targets* in the invisible regions. The positions of the ghost targets are estimated based on the following assumptions: (1) targets tend to stay together as a group, and (2) invisible targets are in C-free outside the visibility region of the camera. Therefore, even if targets are invisible, they must be in some occluded regions nearby. Note that the planner does not distinguish if a target is visible or is a ghost. Therefore the planning strategy described in the previous sections remains the same. Our experiments show that the idea of ghost targets significantly increases the visibility.

5 Discussion and Experimental Results

5.1 Three Additional Following Strategies

We also developed three additional group following strategies. These strategies are extensions of the existing methods which are originally designed to track a single target. Since there is no prior work on group following, we will also use these strategies to compare against the proposed MTR camera.

Reactive Camera. This reactive camera determines its next configuration by placing the visible targets as center in the view as possible based only on the targets' currently positions. The motivation is that by placing the visible targets at the center of its view, the camera will have better chance to make invisible targets visible.

IO Camera. IO camera is a sampling-based method extended from [3]. At each time step, given the visible targets T , the planner first creates k point sets P_T , where k is a parameter. Each point set contains $|T|$ predicted target positions. The predicted position of each target τ is sampled from the region visible

from τ , and is at most $(v_T^{max} \cdot \Delta t)$ away from τ . The planner then creates a set P_C of camera configurations that are at most $(v_C^{max} \cdot \Delta t)$ away from C . To decide the next camera configuration, we simply determine $\arg \max_{x \in P_C} (\sum_{X \in P_T} \text{vis}(x, X))$,

where $\text{vis}(x, X)$ is the number of points in X visible by x . To simplify our discussion, we will use the notation IO- k to denote an IO camera that samples k point sets for target prediction (e.g., in Fig. 3).

VAR Camera. VAR camera is based on [18]. Here, we first obtain a coarse representation of the environment by sampling a grid of discs in C-free. A roadmap is formed over the intersections of the discs. Finally, visibility is computed between each pair of discs with Monte-Carlo raytracing. Our VAR method is a hybrid approach that uses the constructed visibility information in a reactive behavior when the camera has good visibility of the targets (more than 50% of the targets are currently visible by the camera), and uses visibility-aware path planning from [18] to plan short, alternative paths to reach predicted locations of targets when the camera is far away.

The reactive behavior in VAR computes a waypoint for the camera on each frame. First, we find a disc D_c that is closest to the camera, and a disc D_r (from the pre-computed roadmap) that represents an imminent occlusion risk. That is, the disc D_r is one that is in the direction of the visible targets' velocity, closest to the centroid of the visible targets, and whose visibility from the camera's disc is less than 100%. A waypoint is selected along the ray extending from D_r passing through D_c . The visibility-aware path planner is the same as described in [18]. It uses an A* algorithm on the pre-computed roadmap with a heuristic function to compute a path to the D_r which simultaneously minimizes distance traveled and maximizes visibility of the target.

5.2 Experiments and Results

In our experiments, the target group is constantly moving toward a random goal, which is not known by the camera. If all targets are invisible, the camera will stay stationary. Throughout the experiments, we measure the performance of the cameras by computing the *normalized visibility* which is the ratio of the visible targets during the entire simulation. Every data point presented in the plots in this section is an average over 32 runs, each of which is a 10000 time-step simulation. We set the planning horizon $h = 20$ for all MTR cameras.

We perform our experiments in four workspaces (Fig. 3). These workspaces are designed to test the cameras in various conditions, such as large open space (disc), open space with narrow gaps (bars), small irregular obstacles with many narrow gaps (islands) and long sinuous narrow passages (tunnel). Both islands and tunnel environments are considered difficult as the targets tend to separate around the obstacles or hide behind a bend in the passage.

mtr outperforms other strategies. Our first experiment in Fig. 3 shows strong evidence that MTR camera consistently performs better than the other cameras when following 50 targets in all environments. Note that we also include data called *upper bound* obtained from an MTR camera that has no speed

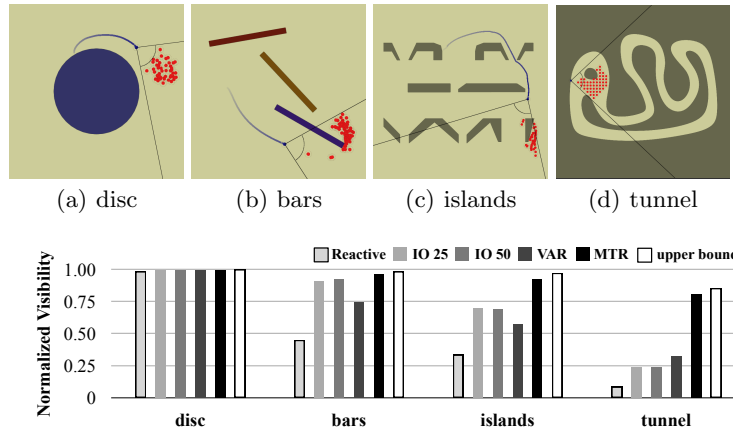


Fig. 3. Following 50 targets in four environments. Normalized visibility is the ratio of the visible targets during the entire simulation. Each bar is an average over 30 runs, each of which is a 9000 time-step simulation.

limitation, i.e., it can move to the best configuration instantly. The strong performance of MTR is further supported by the small difference between the MTR camera and the upper bound in all four environments.

It is clear that the reactive camera performs worst, except in the disc environment. The VAR camera is the second worst, except in the tunnel environment. However, as we will see, the VAR camera seems to handle large and fast moving targets better because of its ability in estimating risks. Although IO cameras perform well in some situations, IO-25 in the bar environment is more than 200 times slower than VAR (≈ 2600 fps) and 12 times slower than MTR (≈ 157 fps), thus cannot be used in many applications, such as real-time task monitoring and video game. There is no significant difference between IO-25 and IO-50.

Ghost targets boost performance. In this experiment shown in Fig. 4, we attempt to estimate quantitatively the performance gain due to the idea of *ghost targets* (GS). Our result shows that the performance gain is more significant in more difficult environments. If compare to the (very time consuming) IO cameras in Fig. 3, MTR without GS is only slightly better in the island and tunnel environments, but there are significant differences between the IO cameras and MTR with GS. In the future, it will be interesting to measure the performance gain by applying GS to reactive and VAR cameras.

6 Conclusion

In this paper, we contribute an online camera planning method, called MTR, which is suitable for autonomous following of multiple targets with unknown trajectories among known obstacles. The idea is to preprocess the known en-




Fig. 4. Comparing MTR cameras with and without ghost target.

environment offline to obtain *monotonic tracking regions* (MTRs) which can then be used to increase the efficiency of the online planner. In addition to MTR, we present three new methods for camera planning that are extensions of existing single-target work, and we compare the performance of MTR with each of these methods. In the tested scenarios, MTR performs significantly better than all other methods, particularly in situations where environments are cluttered with many obstacles or long sinuous tunnels. We also show that adding *ghost targets* to MTR increases the performance.

Acknowledgment

This work is partially supported by NSF IIS-096053 and Autodesk. We also thank the anonymous reviewers for the useful comments.

References

1. Bandyopadhyay, T., Li, Y., Ang Jr., M., Hsu, D.: A greedy strategy for tracking a locally predictable target among obstacles. In: Proc. IEEE Int. Conf. on Robotics & Automation. pp. 2342–2347 (2006)
2. Bares, W.H., Grégoire, J.P., Lester, J.C.: Realtime constraint-based cinematography for complex interactive 3d worlds. In: AAAI '98/IAAI '98. pp. 1101–1106. AAAI (1998)
3. Becker, C., González-Banos, H., Latombe, J.C., Tomasi, C.: An intelligent observer. In: The 4th International Symposium on Experimental Robotics IV. pp. 153–160. Springer-Verlag, London, UK (1997)
4. Bhattacharya, S., Hutchinson, S.: On the existence of nash equilibrium for a two-player pursuit-evasion game with visibility constraints. Int. J. of Rob. Res 57, 251–265 (2009)
5. Butz, A.: Anymation with cathi. In: In Proceedings of the 14th Annual National Conference on Artificial Intelligence (AAAI/IAAI). pp. 957–962 (1997)
6. Courty, N., Marchand, E.: Computer animation: a new application for image-based visual servoing. In: Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation, 2001. vol. 1, pp. 223–228 (2001)
7. Drucker, S.M., Zeltzer, D.: Intelligent camera control in a virtual environment. In: In Proceedings of Graphics Interface '94. pp. 190–199 (1994)

8. Geraerts, R.: Camera planning in virtual environments using the corridor map method. In: in Computer Science (LNCS), S.L.N. (ed.) The Second International Workshop on Motion in Games. pp. 194–209. No. 5884 (2009)
9. Gleicher, M., Witkin, A.: Through-the-lens camera control. In: SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques. pp. 331–340. ACM, New York, NY, USA (1992)
10. Goemans, O., Overmars, M.: Automatic generation of camera motion to track a moving guide. In: In International Workshop on the Algorithmic Foundations of Robotics. pp. 187–202 (2004)
11. Gonzalez-Banos, H., Lee, C.Y., Latombe, J.C.: Real-time combinatorial tracking of a target moving unpredictably among obstacles. In: Proceedings IEEE International Conference on Robotics and Automation. vol. 2, pp. 1683–1690 (2002)
12. Halper, N., Helbing, R., Strothotte, T.: A camera engine for computer games: Managing the trade-off between constraint satisfaction and frame coherence. *Computer Graphics Forum* 20(3), 174 – 183 (2002)
13. Hughes, S., Lewis, M.: Robotic camera control for remote exploration. In: CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems. pp. 511–517. ACM, New York, NY, USA (2004)
14. Jung, B., Sukhatme, G.: A region-based approach for cooperative multi-target tracking in a structured environment. In: Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on. vol. 3, pp. 2764–2769 vol.3 (2002)
15. LaValle, S., Gonzalez-Banos, H., Becker, C., Latombe, J.C.: Motion strategies for maintaining visibility of a moving target. In: Proceedings IEEE International Conference on Robotics and Automation. vol. 1, pp. 731–736 (1997)
16. Li, T.Y., Cheng, C.C.: Real-time camera planning for navigation in virtual environments. In: SG '08: Proceedings of the 9th international symposium on Smart Graphics. pp. 118–129. Springer-Verlag, Berlin, Heidelberg (2008)
17. Murrieta-Cid, R., Tovar, B., Hutchinson, S.: A sampling-based motion planning approach to maintain visibility of unpredictable targets. *Auton. Robots* 19(3), 285–300 (2005)
18. Oskam, T., Sumner, R.W., Thuerey, N., Gross, M.: Visibility transition planning for dynamic camera control. In: SCA '09: Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation. pp. 55–65 (2009)
19. Parker, L.E.: Distributed algorithms for multi-robot observation of multiple moving targets. *Auton. Robots* 12(3), 231–255 (2002)
20. Schulz, D., Burgard, W., Fox, D., Cremers, A.B.: People Tracking with Mobile Robots Using Sample-Based Joint Probabilistic Data Association Filters. *The International Journal of Robotics Research* 22(2), 99–116 (2003)
21. Seligmann, D.D., Feiner, S.: Automated generation of intent-based 3d illustrations. *SIGGRAPH Comput. Graph.* 25(4), 123–132 (1991)
22. Vo, C., Lien, J.M.: Visibility-based strategies for searching and tracking unpredictable coherent targets among known obstacles. In: ICRA 2010 Workshop: Search and Pursuit/Evasion in the Physical World (2010)
23. Wilmarth, S.A., Amato, N.M., Stiller, P.F.: MAPRM: A probabilistic roadmap planner with sampling on the medial axis of the free space. In: Proc. of IEEE Int. Conf. on Robotics and Automation. vol. 2, pp. 1024–1031 (1999)