

# Visibility-Based Strategies for Tracking and Searching Unpredictable Coherent Targets Among Known Obstacles

Christopher Vo

Jyh-Ming Lien

MASC group, Dept. of Computer Science, George Mason University

<http://masc.cs.gmu.edu/wiki/GroupFollowing>

**Abstract**—There are many scenarios in the real world and virtual world where it is necessary to find and track groups of coherent targets, whether they are groups of people, swarms of robots, or players in a video game. In such scenarios, some information about the environment may be available, but trajectories of the targets are usually not known in advance. This paper shows that finding unknown coherent targets in a known environment can be done *asymptotically* more efficiently than finding a single target. In addition, this paper presents several methods for tracking multiple coherent targets with unknown trajectories through environments with known obstacles (such as walls, buildings, etc). We also present variations of our algorithms that incorporate informed search to regain visibility of targets when they escape the camera's view, with promising results.

## I. INTRODUCTION

This paper investigates the problem of *following* and *searching* for a group of coherent targets with unknown trajectories (such as a group of people, a swarm of robots, or a flock of sheep) in a known, obstacle filled environment. We call this problem *group tracking*. The group tracking problem can be found in both the real world and the virtual world. For example, group tracking can be useful in searching for and observing a herd of buffalo at a national park, monitoring swarms of collaborating robots, following an army of Orcs on an obstacle-filled battlefield in Warcraft, or tracking a crowd of protesters marching through a city environment. In these examples, the size of the target group is often large (consisting of tens, hundreds, or perhaps thousands of members) and non-adversarial, the targets usually exhibit some degree of coherence, and the environment is usually known. While there is work on target following and searching (usually for a single target or a small number of targets), to the best of our knowledge, no existing work considers the group tracking problem.

One of the first questions we should ask is: Can we apply existing techniques to solve the group tracking problem? Unfortunately, our preliminary results have found that direct application of existing single-target tracking strategies (see [1] and Table I in Section VII) to track a coherent group usually performs poorly. We speculate that these strategies perform poorly because of the fundamental differences between single-target tracking and group tracking. For example, finding and maintaining the visibility of a coherent group is in some aspects easier than tracking a single target, because the group of targets covers a larger area of the space—one can think of the target group as a single large target. However, keeping

track of a group can be more difficult than a single target if the camera needs to maximize the number of visible targets over time. One of the reasons for this difficulty is that it is possible for a coherent group of targets to change shape and density over time. A group of targets could compress into a long, dense line within a narrow corridor, and then spread out into a sparse blob when it reaches an open area. Groups of targets could also engulf obstacles, or even split into multiple sub-groups for a short period of time. Since the trajectory of the targets is unknown, the trajectory of the camera must also be computed online to react to the movement of the targets. These differences will be discussed in greater depth in Sections IV and VI.

In this paper, we propose methods to both (1) follow and (2) search for a group of coherent targets with unknown trajectories among obstacles. The main ideas of the proposed method are to (1) identify regions with simple (monotonic) topological features so that the planner can repetitively use the same data structure and strategy to follow and search the target group, and (2) utilize the fact the targets form a coherent group. 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 of an 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.

One of the camera's goals is to keep as many targets in the view for as long as possible. In Section IV, we propose four different strategies to do this. These include an approach, called VAR based loosely on [2] which generates a coarse representation of the space that can be used for rapid visibility checking. We make several modifications to the approach. Our method uses a reactive behavior that takes advantage of this data to rapidly compute which areas of the space represent occlusion risk. In Section V, we propose an additional search step to improve the speed of re-acquisition of the target group when it is no longer visible. Finally, in Section VI, we show that finding unknown coherent targets in a known environment can be done *asymptotically* more efficiently than finding a single target.

## II. RELATED WORK

There exist many methods for tracking a target with known trajectory, e.g., [3], [4]. In this section, we will review strategies for tracking a target with unknown trajectory. From our review, we find that no method has focused on the problem of following a large group.

**Tracking Targets with Unknown Trajectories in Known Environment.** There exists some work considering tracking targets with unknown trajectories in a known environment [5], [6], [7]. 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 [8]. Unfortunately, the decomposition usually results in many small components even for a very simple environment. While the nature of workspace decomposition in our MTR method is similar to the combination of [5], [6], MTR creates much fewer components, does not require that the distance between the camera and the target is fixed, and does not have to precompute and store exhaustive visibility information.


Recently, Li and Cheng [9] have proposed a real-time planner that tracks a target with unknown trajectory, using a *budgeted roadmap method* with lazy evaluation and a simple linear extrapolation to predict target’s motion. Oskam et al. [2] developed a *visibility transition planning* method which exhaustively precomputes visibility on a roadmap. This visibility roadmap enables quick prediction of spheres that represent high risk of occlusion so that proactive motion can be taken to prevent this occlusion. The VAR approach we use for comparison in this paper is based on this work.

**Planning with Unknown Target Trajectories in Unknown Environment.** Many researchers have also considered the case where both target trajectory and environment are unknown. For example, Becker et al. [10] proposed a very simple planner to follow a target with unknown trajectory in an environment with landmarks. The idea is to predict the target’s next position and place the camera at the position that can see most of the predicted target positions. Later, González-Baños et al. [11], [12] proposed the main idea of greedily minimizing the *escaping risk* or maximizing the *shortest escaping distance* of the target. Recently, Bandyopadhyay et al. [13], [14], [15] improved the definition of *escaping risk* by introducing the idea of *vantage zone* and showed that this definition can improve the camera’s tracking ability. All of these methods focus on tracking a single target.

**Methods Based on Constraint Satisfaction.** 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., Blinn [16] and Gleicher and Witkin [17]. There are also a number of works which involve the use of metaheuristics. For example, Drucker and Zeltzer [18] have used an A\* planner to compute a camera path. The orientation of the camera is then solved frame-by-frame to satisfy given constraints. Our MTR method is also similar to the idea of *semantic space partitioning*. Christie and Normand [19] used this idea to partition the space so that in each partition, the cinematographic properties remain the same. However, their method is designed to produce a camera view that satisfies certain criteria—not to follow or search for a target.

## III. PRELIMINARIES



We now formally define the problem that we attempt to solve and notation 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 essentially reduces our problem to a 2D workspace. We further assume that during the entire simulation, the set of targets  $T$  exhibit a certain degree of coherence in their motion, and that it is also possible for  $T$  to split

into multiple subgroups for a short period of time (similar to a flock of birds). 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  $\tau$  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  $t$ .

**MTR Definition.** The idea of MTR is used in both search and following. We let a region  $M_\pi$  be a 2D generalized cylinder defined with respect to a *supporting path*  $\pi$ . We say  $\pi$  is a supporting path of  $M_\pi$  if every point  $x \in M_\pi$  can see a subset of  $\pi$ . Because of this property,  $M_\pi$  can essentially be viewed as a linear object and the camera can see every point in  $M_\pi$  by moving along  $\pi$ . Specifically, we define  $M_\pi$  as:  $M_\pi = \{x \mid \exists y \in \pi \text{ s.t. } \overline{xy} \subset C_{free}\}$ , 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  $\pi$  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  $\pi$  if  $|V_\pi(x)| = 1, \forall x \in M_\pi$ , where  $|X|$  is the number of connect components in a set  $X$ . Because each  $x \in M_\pi$

can see only an interval of  $\pi$ , we can compactly represent the visible region (called visibility interval) of  $x$  as a tuple  $V_\pi(x) = (s, t), 0 \leq s \leq t \leq 1$ .

**Target Group.** We consider non-adversarial coherent targets. Many applications fall into this category, e.g., monitoring swarms of robots and following a crowd of video game characters. We model the target group so that the group is constantly moving toward a common goal. When a goal is reached, a new goal (unknown to the camera) is randomly selected, and the group will plan a path toward this new goal. When the target group is moving to a goal, each member in the group keeps itself close to other members (similar to the flocking behavior [20]). More specifically, at every time step, the coherence of the group is maintained by filling the space around a leader in the group. However, when one or multiple obstacles are nearby, the group may separate into multiple sub-groups. In fact, the separation is commonly seen in our experimental setup in Table I in Section VII.

**Tracking.** 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 is then reduced to finding 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, for all  $t$ ,  $v_C(t) \leq v_C^{max}$ , and  $x_C(t)$  is collision free.

**Searching.** When the positions of  $T$  are not known, the camera will be in the search mode. In general, searching is difficult even when the environment is known. In Section VI, we will show that this problem can be significantly simplified when we consider the problem of finding the large crowd.

#### IV. FOLLOWING THE GROUP

We will next discuss four group following strategies. The first three are extensions of existing methods originally designed for tracking a single target. The last strategy is based on the idea of monotonic tracking regions (MTRs).

##### A. Reactive Camera

This reactive camera is our baseline approach. It determines the next configuration by attempting to keep the visible targets centered in its view. Due to the coherence of the group, keeping the visible targets at the center of its view may give the camera a better chance of finding other targets. However, this method does not incorporate search to find the targets.

##### B. IO Camera

IO camera is a sampling-based method extended from [10]. At each time step, given the visible targets  $T$ , the planner first creates  $k$  point sets  $P_T$ , where  $k$  is a user-defined

parameter. Each point set contains  $|T|$  predicted target positions. The predicted position of each target  $\tau$  is sampled from the region visible from  $\tau$ , and is at most  $(v_T^{max} \cdot \Delta t)$  away from  $\tau$ . 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  $\arg \max_{x \in P_C} \left( \sum_{X \in P_T} \text{vis}(x, X) \right)$ , where  $\text{vis}(x, X)$  is the number of points in  $X$  visible by  $x$ . IO cameras usually cannot be used online for tracking large groups (more than 50 targets) because of the high number of visibility checks between the sampled target positions and camera configurations every time step.

##### C. VAR Camera

VAR camera is based on [2]. Here, we first obtain a coarse representation of the environment by sampling a grid of discs  $D$  in C-free. A roadmap is formed over the intersections of the discs. Additionally, visibility is computed between each pair of discs with Monte-Carlo raytracing. We use the constructed visibility information in a reactive behavior when the camera has good visibility of  $T$  (more than 20% of the targets are visible). Otherwise, we use visibility-aware path planning from [2] using A\* and a visibility-based heuristic, to plan short alternative paths to reach predicted locations of targets.

The reactive behavior in VAR computes a waypoint for the camera on each time step. First, we find the disc  $D_c$  that is closest to the camera, and the disc  $D_r$  (from the roadmap) that represents an imminent occlusion risk. That is,  $D_r$  is in the direction of the visible targets' velocity, closest to the centroid of the visible targets, and has visibility from the  $D_c$  less than 100%. A waypoint is selected along the ray extending from  $D_r$  passing through  $D_c$ .

##### D. Monotonic Tracking Region (MTR) Camera

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

MTR camera first finds a set of supporting paths whose MTRs cover  $C_{free}$ , and from a given path  $\pi$ , it computes 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  $\pi$  on the MA forms an MTR (supported by  $\pi$ ).

*Lemma 4.1:* The retraction region  $R \subset C_{free}$  of an edge on the MA forms an MTR. Therefore, the supporting paths are simply constructed by extracting the edges from the MA of a given environment.

**Constructing MTRs.** Given an edge  $\pi$  of MA, its retraction region  $R$  forms an MTR supported by  $\pi$ . However, simply using  $R$  as  $\pi$ 's MTR can be overly conservative. 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. To speed up the computation, our approach is based on the following observation.

*Observation 4.2:* If  $x$  and  $x'$  are (topological) neighbors, and  $x$  is further away from  $\pi$  than  $x'$  is, then  $V_\pi(x) \subset V_\pi(x')$ .

**Following 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  $\tau$ . 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:  $x_C(t + \Delta t) \in V_\pi(x_\tau(t + \Delta t))$ .

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. 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. This can be done by solving an  $h$  dimensional linear programming problem:

$$\begin{aligned} \min \quad & t_h - x_h \\ \text{s.t.} \quad & s_i \leq x_i \leq t_i \\ & 0 \leq (x_{i+1} - x_i) \leq \frac{v_C^{max}}{|\pi|}, \forall x_i, \end{aligned} \quad (2)$$

where  $v_C^{max}/|\pi|$  is the maximum *normalized* distance the camera can travel on  $\pi$ . Finally, the camera's future locations are simply  $x_C(t + \Delta t \cdot i) = \pi(x_i)$ . It is possible that the linear programming problem has no feasible solution. We reduce the plan horizon iteratively until a solution is found.

**Following multiple targets.** Next, we extend this approach to handle multiple targets  $T$ . Let  $x_T(t)$  be the current positions of the targets  $T$ . 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  $I_i = I_i(i \cdot \Delta t) = (s_i, t_i)$ . To maximize the number of visible targets, 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. 2.

**Following the Target between MTRs.** Next, we will discuss strategies to follow the targets across MTRs. Without loss of generality, we only consider the case with two MTRs. The same approach can be naturally extended to handle three or more intersecting MTRs.

Given two MTRs,  $M$  and  $M'$ , we let  $X = M \cap M'$ . The intersection  $X$  of two MTRs plays two critical roles. On one hand, when the target enters  $X$ , the possibility for the target to escape increases. On the other hand, when the camera enters  $X$ , the possibility of seeing the target increases.

Recall that we plan the camera's location by predicting the targets' future locations  $x_T(t + \Delta t)$  and by computing the visible intervals  $I(\Delta t)$  of  $x_T(t + \Delta t)$ . When  $x_T(t + \Delta t)$  reaches  $X$ , each  $x_T(t + \Delta t)$  can have two intervals, one in  $M$  and the other one in  $M'$ . To compute the future locations of the camera, we solve at most *four* canonical following problems. Two of these are in  $M$  and the other two are in  $M'$ . Finally, the camera uses the solution that maximizes the visibility.

## V. REGAINING VISIBILITY OF ESCAPED TARGETS

**VAR camera.** In VAR, we add the following step to help prevent loss of visibility: If less than 20% of  $T$  is visible, the camera computes a prediction of where  $T$  might be next based on  $v_T^{max}$ , and quickly plans a direct path to the predicted location. In a sense, this approach is a "last resort" attempt to catch up with the flock before it completely disappears from sight.

Finally, when the camera has completely lost sight of  $T$ , it can do neither reactive following nor prediction. Our approach in this case is to maintain a "last-viewed timestamp" for each of the discs, and can use this information to prioritize search in least-recently visited areas of the space. The rationale behind this approach is that the timestamp can be updated very quickly across  $D$  due to the availability of a visibility graph.

**MTR camera.** MTR camera 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. 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. 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.

## VI. FINDING THE GROUP

In this section, we briefly show that finding a group of coherent targets with unknown trajectories is easier in general and can be done *asymptotically* more efficiently than finding a single

target. The main idea in our analysis is that by considering the *minimum* area occupied by the group, we can reduce (sometimes drastically) the difficulty of the searching problem, both geometrically and topologically.




Fig. 1. Two cases showing that searching a group of coherent targets can be done asymptotically more efficiently.

Consider the two environments shown in Fig. 1. Each environment presents a scenario in which searching for a group can be done more efficiently than a single target. First, in an indoor-like environment, there can be a long corridor with many rooms (offices or cubicles), which we call *local concavity*. Small local concavity can also appear quite often in modeling errors in workspace geometry (due to numerical, sensor or human errors). In the case of a single target, the camera needs to go into every local concavity in order to ensure that the target is not hiding inside the concavity (e.g.,  $x$  in Fig. 1(a)). Moreover, without multiple collaborative cameras, it may be impossible to exhaustively search the entire workspace in Fig. 1(a) because when the camera enters one of the rooms, other rooms may be (re-)contaminated. However, when we consider a coherent group  $T$  that is large enough to compactly occupy the local concavity such as  $T$  in Fig. 1 (a), the entire workspace can be cleared by simply placing the camera at  $y$  looking toward its left.

Second, in the city-blocks-like environment shown Fig. 1(b), we can see another example where no algorithms can guarantee to find a single target since the environment has many cycles. Therefore, a target can always hide behind an obstacle. In the case of multiple targets, if the group  $T$  is large enough to span over two (horizontal) Homotopy classes (as shown in Fig. 1(b)), the camera can clear the environment by traversing any vertical or horizontal corridor, from one end to the other. The reason why the camera does not need to walk around each obstacle can be shown easily. By placing the camera at  $n$  and then at  $m$ , if the camera cannot see any targets to its right in both places, e.g., at  $n'$  and  $m'$ , (within a certain time) this indicates that there cannot be any targets between  $n'$  and  $m'$  because there is not enough space to fit the entire group  $T$ .

The above mentioned strategies reduce the search complexity to constant and to  $O(n)$  for Figs. 1(a) and (b), respectively, where  $n$  is the complexity of the environment. Both scenarios can be efficiently identified offline since we assume that the environment and the size of the targets are known. More specifically, in every MTR in the environment, we place the group at the end points of its supporting path and determine if the (either local or global) cavities are filled by the group,

TABLE I  
NORMALIZED VISIBILITY FOR ALL FOUR ENVIRONMENTS.

	(a) disc	(b) bars	(c) islands	(d) tunnel
	reactive	IO	VAR	MTR
disc	0.98 (0.03)	0.99 (0.00)	0.99 (0.00)	1.00 (0.00)
bars	0.45 (0.20)	0.91 (0.12)	0.73 (0.13)	0.96 (0.01)
islands	0.34 (0.14)	0.70 (0.17)	0.82 (0.12)	0.93 (0.01)
tunnel	0.09 (0.05)	0.24 (0.12)	0.42 (0.10)	0.81 (0.05)

Each value is averaged over 32 runs. Variances are shown in parentheses.

or if MTR is connected with the neighboring MTRs.

## VII. EXPERIMENTAL RESULTS

In our experiments, the target group is constantly moving toward a random goal that is not known by the camera. 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 of 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, and the IO camera samples 25 point sets for target prediction.

We perform our experiments in four workspaces shown in Table I. These workspaces are designed to test the performance of 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 small obstacles or hide behind obstacles in the passage.

Our first experiment in Table I shows strong evidence that MTR camera consistently performs better than the other cameras when following 50 targets in all environments.

In the experiment shown in Fig. 2, 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. We also compare VAR without the additional search step to VAR with the search step. For this algorithm, adding search to recover visibility makes the most significant difference in the bars environment. We believe that this is because the obstacles are larger and thus it is easier for the entire group to be occluded for long periods of time. While VAR rarely maintains full view of the flock in the tunnel environment, it also rarely completely loses the flock. Therefore, search is not as relevant in the tunnel environment.




Fig. 2. Comparing MTR cameras with and without ghost target (top) and Comparing VAR with and without search step (bottom)

## VIII. CONCLUSION AND FUTURE WORK

This paper contributes online camera planning methods that can track multiple targets with unknown trajectories in an environment with known obstacles. Our first method uses offline computation on the known environment to obtain *monotonic tracking regions* (MTRs) which can be used for more efficiently searching the space and following targets by taking. We also show an extension that samples predictive "ghost targets" that further improves the efficiency of the approach by taking advantage of target coherence. For comparison, we also present methods called IO and VAR which are loosely based on an existing single-target tracking techniques. In particular, with VAR, we show that adding an additional informed search step can greatly enhance the performance of the approach in cases where the camera loses sight of the targets.

We believe that in some situations, multiple cameras may be needed to achieve better performance. Therefore, we are working on applying this method to tracking scenarios involving multiple cameras, and scenarios involving multiple objectives (such as quality, or cinematographic preferences).

## ACKNOWLEDGMENT

This work is supported in part by NSF grant IIS-096053 and by Autodesk. We also thank the anonymous reviewer for the useful comments.

## REFERENCES

- [1] C. Vo and J.-M. Lien, "Following a large unpredictable group of targets among obstacles," George Mason University, Tech. Rep. GMU-CS-TR-2010-3, January 2010.
- [2] T. Oskam, R. W. Sumner, N. Thuerey, and M. Gross, "Visibility transition planning for dynamic camera control," in *SCA '09: Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. New York, NY, USA: ACM, 2009, pp. 55–65.

- [3] S. LaValle, H. Gonzalez-Banos, C. Becker, and J.-C. Latombe, "Motion strategies for maintaining visibility of a moving target," in *Robotics and Automation, 1997. Proceedings., 1997 IEEE International Conference on*, vol. 1, Apr 1997, pp. 731–736 vol.1.
- [4] O. Goemans and M. Overmars, "Automatic generation of camera motion to track a moving guide," in *In International Workshop on the Algorithmic Foundations of Robotics*, 2004, pp. 187–202.
- [5] R. Murrieta-Cid, B. Tovar, and S. Hutchinson, "A sampling-based motion planning approach to maintain visibility of unpredictable targets," *Auton. Robots*, vol. 19, no. 3, pp. 285–300, 2005.
- [6] R. Murrieta-Cid, T. Muppurala, A. Sarmiento, S. Bhattacharya, and S. Hutchinson, "Surveillance Strategies for a Pursuer with Finite Sensor Range," *The International Journal of Robotics Research*, vol. 26, no. 3, pp. 233–253, 2007.
- [7] S. Bhattacharya, S. Candido, and S. Hutchinson, "Motion strategies for surveillance," in *Proceedings of Robotics: Science and Systems*, Atlanta, GA, USA, June 2007.
- [8] S. Bhattacharya and S. Hutchinson, "On the existence of nash equilibrium for a two-player pursuit-evasion game with visibility constraints," *Int. J. of Rob. Res.*, vol. 57, pp. 251–265, 2009.
- [9] T.-Y. Li and C.-C. Cheng, "Real-time camera planning for navigation in virtual environments," in *SG '08: Proceedings of the 9th international symposium on Smart Graphics*. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 118–129.
- [10] C. Becker, H. González-Banos, J.-C. Latombe, and C. Tomasi, "An intelligent observer," in *The 4th International Symposium on Experimental Robotics IV*. London, UK: Springer-Verlag, 1997, pp. 153–160.
- [11] C.-Y. Lee, H. González-Baños, and J.-C. Latombe, "Real-time tracking of an unpredictable target amidst unknown obstacles," 2002.
- [12] H. Gonzalez-Banos, C.-Y. Lee, and J.-C. Latombe, "Real-time combinatorial tracking of a target moving unpredictably among obstacles," in *Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on*, vol. 2, 2002, pp. 1683–1690 vol.2.
- [13] T. Bandyopadhyay, Y. Li, M. H. A. A. Jr, and D. Hsu, "Stealth tracking of an unpredictable target among obstacles," *Algorithmic Foundations of Robotics VI*, vol. 17, pp. 43–58, October 2005.
- [14] T. Bandyopadhyay, Y. Li, M. Ang Jr., and D. Hsu, "A greedy strategy for tracking a locally predictable target among obstacles," in *Proc. IEEE Int. Conf. on Robotics & Automation*, 2006, pp. 2342–2347.
- [15] T. Bandyopadhyay, M. Ang Jr., and D. Hsu, "Motion planning for 3-d target tracking among obstacles," in *Proc. Int. Symp. on Robotics Research (ISRR)*, 2007.
- [16] J. Blinn, "Where am I? What am I looking at?" *IEEE Computer Graphics and Applications*, vol. 8, no. 4, pp. 76 – 81, July 1988.
- [17] M. Gleicher and A. Witkin, "Through-the-lens camera control," in *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM, 1992, pp. 331–340.
- [18] S. M. Drucker and D. Zeltzer, "Intelligent camera control in a virtual environment," in *In Proceedings of Graphics Interface '94*, 1994, pp. 190–199.
- [19] M. Christie and J.-M. Normand, "A semantic space partitioning approach to virtual camera composition," *Computer Graphics Forum*, vol. 24, no. 3, pp. 247 – 256, 2005.
- [20] C. W. Reynolds, "Flocks, herds, and schools: A distributed behavioral model," in *Computer Graphics (SIGGRAPH '87 Proceedings)*, vol. 21, Jul. 1987, pp. 25–34.