Dealing with geometric constraints in game-theoretic planning

Fabiani, Patrick
Visiting Scholar at Stanford
ONERA/ DCSD-Toulouse,
2 av. Edouard Belin, BP4025,
Toulouse, F-31055, Prance.

Latombe, Jean-Claude
Stanford University
Department of Computer Science
Robotics Laboratory
Stanford, CA 94305-9010, USA.

Abstract

Tracking a partially predictable moving object in a workspace cluttered by obstacles is a challenging combination of two planning problems. The planner must take into account on the one hand the visibility and motion constraints imposed by the obstacles and on the other hand the uncertainties in both the robot's position and the future trajectory of the target. Game theory is proposed as an appropriate framework to solve this twofold problem. At each time step, a probability distribution models the positioning uncertainties of the robot and the target and a utility function represents the reward associated with the possible goal states of the motion decision problem. This approach allows the simple modeling of different tracking strategies, one of which has been implemented and tested successfully with two mobile robots. By considering simultaneously target visibility and position uncertainty, the tracking robot can take advantage of landmarks scattered in the workspace to better localize itself so as to better track the target afterward. Future extensions seem possible in order to introduce useful symbolic inferences within the game-theoretic framework.

1 Introduction

Dealing with uncertainty is a major problem in Robotics (e.g. see [Dorst *et* al., 1996] for an overview), which raises and combines a number of issues often associated with estimation, planning or execution control. It is usually unrealistic to assume that a complete and absolutely precise model of the environment can ever be obtained. Yet, a robot is expected to compute and execute robust plans. While Computational Geometry offers tools to deal with geometric constraints (collision, visibility) arising from the robot's workspace, Artificial Intelligence techniques make it possible to make decision in the presence uncertainty. But the techniques from those two fields have rarely been combined and in fact are seldom compatible.

In this paper we show how methods developed independently can be combined to solve a target tracking problem presented in § 2. In this problem, a mobile robot must keep a moving target into its field of view. The workspace is cluttered by obstacles that create motion and visibility constraints. Uncertainties in the robot's self-localization and in the target behavior add to the complexity of the problem. Game theory [Luce and Raiffa, 1957] offers a convenient framework for both decision making under uncertainty [Savage, 1972] and robot motion planning [LaValle, 1995]. The game-theoretic formulation presented in § 4 is specifically proposed for the online computation and execution of target-tracking motion strategies in the presence of uncertainties (§ 6). We have implemented the techniques presented in this paper and tested them with real robots: characteristics of the produced motion strategies are discussed in § 7. Some shortcomings lead us to propose several extensions in § 8.

2 Tracking problem

A mobile robot equipped with visual sensors, called the *observer,* operates in a workspace cluttered with static obstacles. Its task is to keep another moving object, the *target,* in its field of view while avoiding collision with the obstacles. All the obstacles create visual occlusion. Hence, the observer must move in such a way that the target is never hidden behind an obstacle.

We assume that the workspace is realistically described by a 2-D map (planar layout) of the obstacles. Both the observer and the target are modeled as discs. An accurate map of the obstacles is given to the observer before target tracking begins.

The observer uses a camera mounted on a turret to detect the target. The turret rotates fast enough to make it possible to consider that the observer has omnidirectional vision, independent of its current heading. The observer detects and recognizes the target whenever it has a cone of sight to the entire target's disk not obstructed by any obstacle. In our implementation, a black-and-white geometric pattern is mounted on the target to simplify visual detection. There exist more sophisticated vision techniques to detect natural targets (e.g. see [Bregler and Malik, 1998]).

The target is mostly unpredictable, that is, its future trajectory is not known in advance, but it moves collision-free and its velocity has a known upper bound. The observer localizes itself relative to the workspace using both odometrie and landmark techniques. Our *landmarks* are simple geometric patterns placed at the ceiling of the robot environment; they are detected and recognized using a camera mounted vertically on the observer platform [Becker *et* al., 1995]. Natural landmarks could be used as well, but would require more sophisticated vision techniques. Whenever the observer sees a landmark, it localizes itself in the workspace with a given precision (on the order of l-2in) that has been experimentally established. When the robot does not see any landmark, the imprecision of its location estimate (dead-reckoning) grows linearly with the distance traveled. To keep our work realistic and compatible with the use of natural landmarks, the artificial landmarks are sparsely distributed over the workspace. The two vision systems used for target tracking and landmark detection, respectively, are independent and operate concurrently.

The observer uses the vision system that tracks the target to estimate the target's position relative to itself and (through a simple transform) to the workspace. Note that the observer cannot track the target by only using their position relative to one another. It must also estimate both its own position relative to the workspace to avoid colliding with the obstacles and the target's position in the workspace in order to move in such a way that it prevents the target from hiding behind an obstacle.

3 Related work

This tracking problem is a challenging combination of different planning problems that have previously been studied in Robotics, but separately. Visibility and collision constraints must be satisfied in the presence of uncertainties in the positions of the observer and the target. The problem of maintaining visibility with the target while avoiding collision with the obstacles is addressed in [LaValle *et* al., 1997], but position uncertainties are not taken into account. When the target is fully predictable, that is, when its future trajectory is completely known in advance, a dynamic programming approach [Bertsekas, 1986] can be used to compute a trajectory of the observer that has minimal length.

As shown in [LaValle *et a*/., 1997], this approach becomes untractable in practice if the target is only partially predictable. Then, the approach can be applied to choose a motion command that aims to maximize the likelihood that the target will remain visible during a short interval of time in the future. The observer iterates this computation while tracking the target and updates its motion heading at each iteration. This is essentially the approach taken in this paper. The main difference is that we take into account the uncertainties in the observer's and target's positions. This difference will motivate the embedding of this approach into the game-theoretic framework presented in § 4.

Taking position uncertainties into account naturally leads to planning the observer's motions in order to take advantage of the landmarks in the workspace to reduce the imprecision of the estimate of the observer's position (and consequently that of the target), whenever this does not immediately conflict with keeping the target in the observer's field of view. Landmark-based navigation has been addressed from different points of view in the literature (e.g. see [Bouilly and Simeon, 1996; Saffiotti and Wesley, 1996; Takeda *et* al., 1994; Lazanas and Latombe, 1995]). The principle is simple: if the robot primarily localizes itself relative to landmarks, the planner must guarantee that the robot will see landmarks often enough along its path [Fualdes and Barrouil, 1993]. To our knowledge, no previous work on landmark-based navigation simultaneously deals with achieving a visual task such as maintaining visibility with a moving target.

4 Game-theoretic approach

Different discrete processes intervene in the behavior of the observer. During each iteration, the observer chooses its next move according to the last estimate of the target's position, executes it, localizes itself (by using a landmark, if one is visible), and computes a new estimate of the target's position. Meanwhile, the target is moving to another location.

Let t_1 be the initial time point, Δt the duration of the time interval between two sampling time (the duration of the global decision loop), and $l \geq 1$ an index corresponding to the time point $t_k = t_1 + (k-1) * \Delta t$. The observer and the target move in the same workspace, but since their respective shapes and sizes may differ, their free configurations belong to two different configuration spaces, defined with standard parameterizations [Latombe, 1991], respectively C_{free}^{obs} and C_{free}^{tgt}. Let q_k^{obs} and q_k^{tgt} denote the respective configurations of the *observer* and the *target* at time *k*. Then, the global state at time *k* is given by the couple $\omega_k = (q_k^{obs}, q_k^{tgt}) \in C_{free}^{obs} \times C_{free}^{tgt}$. Let $\Omega = C_{free}^{obs} \times C_{free}^{tgt}$ denote the set of all possible states. Let $\mathcal{V}^{tgt}(q^{tgt})$ designate the *visibility region* of the target at configuration q^{tgt}, i.e., the set of all configurations q^{obs} from which the observer has an unobstructed cone of sight to the entire body of the target. Let $\mathcal{R}_{k+1}^{obs}(q_k^{obs})$ (resp. $\mathcal{R}_{k+1}^{tgt}(q_k^{tgt})$) stand for the set of all the configurations in C_{free}^{obs} (resp. C_{free}^{tgt}) that are reachable without collision q_k^{obs} (resp. q_k^{tgt}) between times k and $k+1$.

Game theory provides a convenient framework to express decision problems in the presence of uncertainty. Yet, most work in that field has focused on the probabilistic part of the theory, especially on the Bayesian computation of the probability distributions (POMDP [White, 1993], Bayesian Networks [Pearl, 1988]). Game theory rests on the axiomatization of a relation of preference defined over the possible choices. This relation rests on the definition of both a probability measure and a utility function over the possible eonsequences. We think that the utility part of the theory is as important as the probabilistic part.

Let us define a decision problem over the triplet $\{\Omega, \mathcal{D}, \mathcal{K}\}$, where Ω is the set of atomic *states,* $\mathcal{D}$ is the set of possible *decisions,* and $\mathcal{K}$ is the set of possible *consequences* $k_{ij} = d_j(\omega_i)$. Each possible decision d_j is a function $\{d_j : \Omega \longrightarrow \mathcal{K}\}$.

A decision problem also requires the definition of a *relation of preference* $\succeq$ over $\mathcal{D} \times \mathcal{D}$. For each $(d, d') \in \mathcal{D} \times \mathcal{D}$, $d \succeq d'$ must be read "*d* is preferable than d'" Von Neumann's theorem tells us that this relation of preference can be defined as follows [Von Neumann and Morgenstern, 1944]:

- Let *P* be a *probability distribution* over the set Ω of states,
- Let *U* be a *utility function* over the set $\mathcal{K}$ of consequences,
- Posing:

$$\mathrm{E}_P U(d) = \sum_{\omega \in \Omega} P(\omega).U(d(\omega)),$$

$(\mathrm{E}_P U(d)$ is the expected utility of choosing $d \in \mathcal{D})$ we have:

$$d \succeq d' \quad \text{if and only if} \quad \mathrm{E}_P U(d) \geq \mathrm{E}_P U(d').$$

With this formulation, the best decision d^* is:

$$d^* = arg \max_{d \in \mathcal{D}} \{\mathrm{E}_P U(d)\} \tag{1}$$

In our target tracking problem, suppose that at time k+1 the target is at configuration q^{tgt}_{k+1} and the observer has reached q^{obs}_{k+1} after the motion decision d^{obs}_k was made in state (q^{obs}_k, q^{tgt}_k). The new state $(q^{obs}_{k+1}, q^{tgt}_{k+1})$ is a consequence of dk. The need that q^{obs}_{k+1} be inside the visibility region $\mathcal{V}^{tgt}(q^{tgt}_{k+1})$ can be expressed in the utility function of the problem.

More precisely, our tracking problem is a sequence of motion decision problems (one at each time k) that are all defined on the same state space $\Omega = \mathcal{C}^{obs}_{free} \times \mathcal{C}^{tgt}_{free}$. The set $\mathcal{K}$ of possible *consequences* is also fixed. As the observer may collide with obstacles by mistake (due to localization errors), there is an additional consequence *ceil* (standing for "collision"). The resulting set of consequences is $\mathcal{K} = \{\mathcal{C}^{obs}_{free} \cup \{coll\}\} \times \mathcal{C}^{tgt}_{free}$. At each time k the motion decision problem can be defined as follows:

- The observer's and target's positions are not precisely known at time k. Let P^{obs}_k (resp. P^{tgt}_k) denote the probability distribution of the observer's (resp. target's) position over $\mathcal{C}^{obs}_{free}$ (resp. $\mathcal{C}^{tgt}_{free}$). The joint distribution $P = (P^{obs}_k, P^{tgt}_k)$ is the probability distribution over Ω.
- Let $\tilde{q}^{obs}_k$ be the current estimate of the position of the observer according to P^{obs}_k. This estimate can be a singleton or a set of configurations.
- Let $\mathcal{D}^{obs}_k$ be the set of possible motion decisions for the observer from the current estimated configuration $\tilde{q}^{obs}_k$. In practice, $\mathcal{D}^{obs}_k$ is finite.
- The utility function U_k over $\mathcal{K}$ (here, the set of possible consequences at time k 4- 1) can be simply defined to be 1 if the observer sees the target at time k + 1 and 0 if it does not see the target or has collided with an obstacle. More formally:

$$U_k(coll, q^{tgt}_{k+1}) = 0 \tag{2}$$

$$U_k(q^{obs}_{k+1}, q^{tgt}_{k+1}) = 1 \quad \text{if} \quad q^{obs}_{k+1} \in \mathcal{V}^{tgt}(q^{tgt}_{k+1}) \tag{3}$$

$$= 0 \quad \text{otherwise} \tag{4}$$

Figure 1: Geometric constraints and utility regions

Geometric constraints illustrated in Fig. 1 directly affect the evaluation of the above utility function: Is the observer colliding with obstacles, or not? Is the observer in the visibility region of the target? As we will see in the next section, they also affect the computation of the probability distributions P_k and P_{k+1}, hence that of the best motion decision defined by Equ. (1): can the observer see a landmark?

Note that more complex utility functions could be considered. For example, $U_k(q^{obs}_{k+1}, q^{tgt}_{k+1})$ could evaluate to the minimal distance that the target would have to travel before q^{obs}_{k+1} is no longer in the target's visibility region.

5 Modeling uncertainties

In our work will only consider uncertainties caused by inaccuracies in sensing in motion control and in lack of prior knowledge of the target's trajectory. Uncertainty in the observer's location depends critically on whether it sees a landmark, or not.

Whenever the observer moves and no landmark is visible, dead-reckoning errors occur and accumulate. Consequently, the uncertainty on the position of the observer increases (see Fig. 2) and the observer may risk collision along a supposedly collision-free path. The actually reached configuration at time k- 1, q^{obs}_{k+1}, can be modeled as a random variable that is conditionally dependent upon both the observer's configuration at time k, q^{obs}_k, and the decision made at time k, d^{obs}_k.

The conditional probability $Pr(q^{obs}_{k+1} \mid d^{obs}_k, q^{obs}_k)$ is given *a priori* and allows the computation (at time *k)* of the predictive probability $\hat{P}^{obs}_{k+1}$ law at time $k+1$ which models the imprecision of dead-reckoning.

$$\forall q^{obs} \in \mathcal{C}^{obs}_{free},\ \hat{P}^{obs}_{k+1}(q^{obs} \mid d^{obs}_k) = \sum_{q^{obs}_- \in \mathcal{C}^{obs}_{free},\, d^{obs}_k \in \mathcal{D}^{obs}_k} Pr(q^{obs} \mid q^{obs}_-, d^{obs}_k)\, P^{obs}_k(q^{obs}_-) \quad (5)$$

$$\hat{P}^{obs}_{k+1}(coll \mid d^{obs}_k) = 1 - \sum_{q^{obs} \in \mathcal{C}^{obs}_{free}} \hat{P}^{obs}_{k+1}(q^{obs} \mid d^{obs}_k) \quad (6)$$

If localization fails at time $k+1$, the predictive probability law $\hat{P}^{obs}_{k+1}$ cannot be updated: then the best estimate of the observer's location at time $k+1$ is the one given by dead-reckoning and $P^{obs}_{k+1} = \hat{P}^{obs}_{k+1}$ (the predictive probability law). The location uncertainty reduces whenever

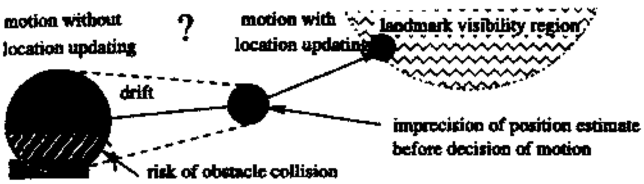

Figure 2: The localization dilemma

the observer sees a landmark. Let loc^{obs}_{k+1} be a boolean variable equal to 1 whenever the observer reaches a position at time k -f 1 where it can see a landmark, and 0 otherwise (or if the observation of the landmark fails). When a landmark is visible, the positioning accuracy that is obtained after the updating process at time $k+1$ depends both on the predictive probability law $\hat{P}^{obs}_{k+1}$, and on the experimentally established precision of the localization. This is modeled by the updating function $\mathcal{L}^{obs}$ in the following equation:

$$P^{obs}_{k+1} = \mathcal{L}^{obs}(\hat{P}^{obs}_{k+1}) \quad \text{if and only if } loc^{obs}_{k+1} = 1 \quad (7)$$

$$= \hat{P}^{obs}_{k+1} \quad \text{if and only if } loc^{obs}_{k+1} = 0 \quad (8)$$

Eventually, the probability distribution P^{obs}_{k+1} at time k depends both on $\hat{P}^{obs}_{k+1}$ and on the observation of a landmark, if a landmark has been observed. P^{obs}_{k+1} is thus given by equations (5) to (8) as a function of the probability measure P^{obs}_k, the decision d^{obs}_k made at time k, and the value of the variable loc^{obs}_{k+1}. This is summarized in the following *evolution equation,* where the *evolution function* $\mathcal{F}^{obs}$ models both the possible motion errors and the reduction of positioning imprecision whenever a *landmark* is detected:

$$P^{obs}_{k+1} = \mathcal{F}^{obs}(P^{obs}_k, d^{obs}_k, loc^{obs}_{k+1}) \quad 0)$$

$$P^{obs}_{k+1}(coll) = 1 - \sum_{q^{obs} \in \mathcal{C}^{obs}_{free}} P^{obs}_{k+1}(q^{obs}) \quad (10)$$

Similarly, due to the limited capabilities of its tracking camera, the observer can only obtain an imprecise estimate of the position of the target relative to own current position (see Fig. 3). Let $\tilde{P}^{tgt}_k$ denote the probability over the possible configurations of the target after detection by the observer at time k. The uncertainty on the position of the target in the workspace combines the uncertainty on the target's position relative to the observer and the uncertainty on the observer's position in the workspace. Hence, $\tilde{P}^{tgt}_k$ depends on both P^{tgt}_k and P^{obs}_k and it is necessarily more imprecise than P^{obs}_k (see Fig. 3). The target's partial predictability between times k and $k+1$ can be written: $P^{tgt}_{k+1} = \mathcal{S}^{tgt}(\tilde{P}^{tgt}_k)$. Since the target's trajectory is not known in advance, S^{tgt} depends only the known target's maximal velocity. This means that the uncertainty on the target's position at time $k+1$ (as predicted at time k) may be much greater than for the observer. On the other hand, the target is observed at every time step before the prediction stage, which considerably reduces uncertainty. The evolution equation for the target takes into account how the probability measure P^{tgt}_k is updated at time k by the observer and how it evolves from time k to time k -f 1 due to partial predictability of the target:

$$P^{tgt}_{k+1} = \mathcal{F}^{tgt}(P^{tgt}_k, P^{obs}_k) \quad (11)$$

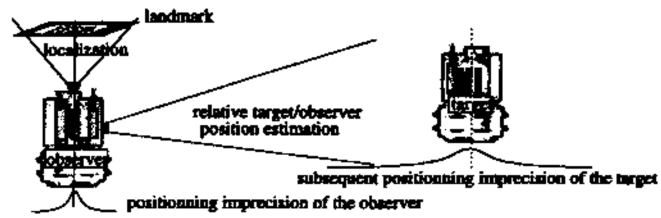

Figure 3: Positioning imprecision

6 Tracking strategies

At each time step, the observer, decides where to move according to P^{tgt}_k and P^{obs}_k. $P_k = (P^{obs}_k, P^{tgt}_k)$ is the *information state* of the observer at time k and the global *evolution equation* $P_{k+1} = \mathcal{F}(P_k, d^{obs}_k, loc^{obs}_{k+1})$ follows from Equ. (9) and (11). A strategy π for the tracking problem with finite time horizon $\{1, h\}$ [White, 1993] is a sequence $\{\ \pi_k\ \}_{k \in \{1,h\}}$ of *decision rules* $\pi_k(P_k) = d^{obs}_k$. An optimal strategy is obtained by discrete dynamic programming [White, 1993]. Let denote $V^h(P_1)$ the *optimal value function* for the /i-stage tracking problem beginning with an initial probability distribution P_1 on Ω at time *t1*, i.e. the maximum expected reward that the observer can get with this initial condition. Bellman's optimality equations give : $\forall k \in \{1, h\}$

$$V^{h-k+1}(P_k) = \max_{d \in \mathcal{D}^{obs}(P_k)} \left\{ \mathrm{E}_{P_{k+1}} U_k(d) + V^{h-k}(P_{k+1}) \right\} \quad (12)$$

where $V^0(.) = 0$ and $P_{k+1} = \mathcal{F}(P_k, d, loc^{obs})$.

Unfortunately, as already mentioned in § 3, the cost of computing an optimal strategy is exponential in the horizon *h* (number of time steps). This difficulty is well-known in the field of POMDP's [White, 1993] where dynamic programming is nevertheless widely used. In target tracking, a greater h would make it possible to compute better strategies if the computation time was negligible. But since this is not the case and the target moves while the observer computes its next motion, a compromise is necessary. In our experiments in simulation, we computed strategies with h = 2. In our experiments with real robots, we used h = 1. The empirically established real-time constraint with real robots is that the observer must choose a motion command within 50ms. On a SUN SPARCstation 20 the computation time for ft = 1 ranges between 1 and 10ms, but grows up to 120-1200ms for h = 2.

For h = 1, the decision rule at time k when P_k is the probability distribution over Ω corresponds to $V^1(P_k)$ in Equ. (12). Hence, the best decision is:

$$d_k^*(P_k) = arg \max_{d \in \mathcal{D}^{obs}(P_k)} \left\{ \mathrm{E}_{P_{k+1}} U_k(d) \right\} \qquad (13)$$

where $P_{k+1} = \mathcal{F}(P_k, d, loc^{obs})$.

7 Experimentation

Experiments were conducted to test our game-theoretic approach : the observer and the target are two Nomad-200 mobile robots of Nomadic Technologies. Each Nomad-200 has a rigid cylindrical body that allows its representation as a disc in the 2-D workspace. The *observer* detects landmarks placed at the ceiling of the laboratory. It can track the target regardless of its current heading, thanks to a camera that is kept aimed at the target by servoing the observer's turret appropriately. All the obstacles in the environment are polygonal and obstruct both visibility and motion. More sophisticated geometric models for the obstacles are possible. Simple probabilistic models of the robots were introduced in the evolution equations (Equ. (9)-(II)): they can be improved experimentally. Their good behavior in practice is due to the fact that they are rather conservative. Yet, it can be argued that little improvement is to be expected from more precise probabilistic models, while so little is known about the target's future motion.

Thanks to the fact that all obstacles are static and known in advance, a number of precomputations are performed in advance in order to reduce on-line computation times (e.g. visibility and reachability regions). The utility function U_k at each time step k is obtained through a small amount of calculation (Equ. (2)-(4)) and with nice properties of piecewise smoothness. U_k could have been chosen to give a higher penalty to collision. For a given state (q_k^{obs}, q_k^{tgt}) at time k, a more sophisticated U_k could have been proposed, for example to penalize states where the observer is close to loosing the target and/or where the observer's turret needs to rotate by a large amount. U_k could also be designed to take into account the speed of the observer at stage k in order to advantage future configurations that correspond to small accelerations of the *observer* in either direction. Yet, simplicity was preferred.

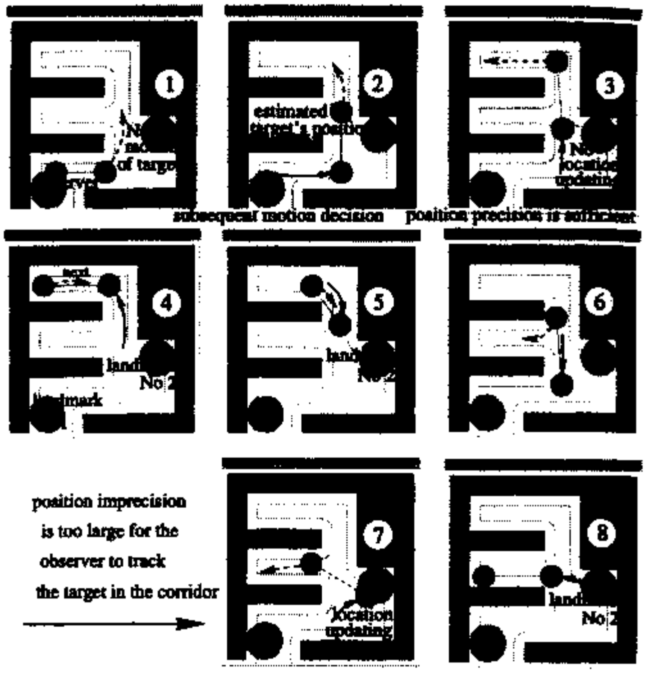

Figure 4: Sample sequence of motion decisions

Fig. 4 represents an 8-step decision sequence generated by the observer's planner in response to target's moves. This sample sequence shows snapshots of the graphic display of the conducted experiments: the target appears as a dark disk and the observer as a light-grey disc with a little tail. In each snapshot, the small grey circle with variable radius and the same center as the observer depicts the observer's positioning uncertainty. The other static light-grey disks represent landmark regions (regions in which the observer's center must lie in order to see a landmark). The landmark that was detected the most recently, recognized and used for location updating by the *observer* is marked with a small blue circle in the center. The arrows describe each agent's moves. Please note that the *target's* moves are represented by arrows reflecting the *target's next move* in the scenario, whereas for the *observer* the arrows reflect the *subsequent move* that is performed by the *observer* in response to the *observed move* of the target at the *previous stage.*

The conducted implementation successfully validates our approach. The *observer* is able to manage properly its positioning uncertainty. Whenever a *landmark visibility* region is reachable, the motion decision corresponds to the *best compromise* between the tracking task and the need for location updating. There is no conflict because the *observer* localizes itself so as to better track the target afterward. More precisely, the following features are remarkable in its behavior in Fig. 4:

- the *observer* takes its positioning uncertainty into account before choosing its next move: it localizes itself when useful, but does not need to do so when it can travel in a wide enough free area (stage 1-6);

- the *observer* goes and localizes itself so as to reduce its positioning imprecision (if possible) prior to entering a narrow region (stages 7-8);
- the *observer* always minimizes its probability of collision because this improves the probability of visibility of the *target* at next time step;
- in practice the *observer's* probability of collision remains equal to zero whenever possible (stages 1-8 here), unless the *observer* gets trapped by the *target* in a narrow region where it cannot reach any *landmark* region in order to localize itself.

8 Perspectives

The nice experimental behavior of the designed planner shows that the choice of modeling this problem in a game-theoretic framework was more than a theoretical exercise: it should allow geometric algorithms to be efficiently combined with algorithms designed for POMDP's or Bayesian Networks. Still, this would need to be further formalized with more generality, so to be applicable to a wider range of problems.

Furthermore, the proposed game-theoretic framework allows dealing with symbolic constraints (e.g. *"a robot cannot see/travel through walls")* and reasoning with probabilistic uncertainty to be efficiently combined through the computation of well-adapted utility functions. The introduction of reasoning mechanisms, such as in Fuzzy Sets theory [Saffiotti and Wesley, 1996] or Possibility theory [Dubois and Prade, 1995] seem to be a promising approach for reasoning with utility functions.

In that perspective, the conducted experiments provide interesting ideas for future developments. For instance, when the *observer* happens to estimate the position of the *target* on a configuration that is not collision-free, this may be due to location and detection errors, but also to the fact that some obstacles in the modeled workspace of the robots do not exist in the real environment: our 2-D layout of the workspace actually is a very conservative model of the real indoor environment. Dealing with an incomplete and inaccurate map in full generality is a really tough challenge. The proposed generalized game-theoretic approach should provide an interesting framework in order to release progressively the hypothesis of a perfectly known environment.

Acknowledgments

We thank Hector H. Gonzales-Banos and David Lin for their technical inputs throughout this research and for their assistance in setting up the robotic experiments. This research was funded in part by ARO MURI grant DAAH04-96-1-007. Patrick Fabiani was supported by DGA and ONERA.

References

[Becker *et* al., 1995] C. Becker, J. Salas, K. Tokusei, and J.C. Latombe. Reliable navigation using landmarks. In *ICRA,* 1995.

[Bertsekas, 1986] D.P, Bertsekas. *Dynamic programming : deterministic and stochastic models.* Prentice-Hall, 1986.

[Bouilly and Simeon, 1996] B. Bouilly and T. Simeon. A sensor-based motion planner for mobile robot navigation with uncertainty. In Dorst et al. [1996].

[Bregler and Malik, 1998] C. Bregler and J. Malik. Tracking people with twists and exponential maps. In *IEEE Computer Society Conf. on Computer Vision and Pattern Recognition,* pages 8-15, 1998.

[Dorst *et al.,* 1996] L. Dorst, M. van Lambalgen, and F. Voorbraak, editors. *Reasoning with Uncertainty in Robotics.* Springer, mar 1996.

[Dubois and Prade, 1995] D. Dubois and H. Prade. Possibility theory as a basis for qualitative decision theory. In *IJCAI95,* pages 1924-1930. Montreal, 1995.

[Fualdes and Barrouil, 1993] T. Fualdes and C. Barrouil. Perception planning for an U.G.V. *Robotics and Autonomous Systems,* 11(2), 1993.

[Latombe, 1991] J.C. Latombe. *Robot Motion Planning.* Kluwer Academic Publishers, Boston, 1991.

[LaValle *et al.,* 1997] S.M. LaValle, H. Gonzaiez-Banos, C. Becker, and J.C. Latombe. Motion strategies for maintaining visibility of a moving target. In *ICRA,* 1997.

[LaValle, 1995] S.M. LaValle. *A Game-Theoretic Framework for Robot Motion Planning.* PhD thesis, University of Illinois, Beckman Institute, 1995.

[Lazanas and Latombe, 1995] A. Lazanas and J.C. Latombe. Motion planning with uncertainty: a landmark approach. *Artificial Intelligence,* 76(1-2), 1995.

[Luce and Raiffa, 1957] R. D. Luce and H. Raiffa. *Games and Decisions.* John Wiley & Sons, 1957.

[Pearl, 1988] J. Pearl. *Probabilistic reasoning in intelligent systems : networks of plausible inference.* Morgan Kaufmann, San Mateo, 1988.

[Saffiotti and Wesley, 1996] A. Saffiotti and L.P. Wesley. Perception-based self-localization using fuzzy locations. In Dorst et al. [1996], pages 368-385.

[Savage, 1972] L.J. Savage. *The Foundations of Statistics.* Dover, New York, 1972.

[Takeda *et al,* 1994] H. Takeda, C. Facchinetti, and J.C. Latombe. Planning the motions of a mobile robot in a sensory uncertainty field. *IEEE Trans, on Pattern Analysis and Machine Intelligence,* 16(10), 1994.

[Von Neumann and Morgenstern, 1944] J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, 1944.

[White, 1993] D.J. White. *Markov Decision Processes.* John Wiley & Sons, Chichester, 1993.