

Active Localization of People with a Mobile Robot Based on Learned Motion Behaviors

Maren Bennewitz Julio Pastrana Wolfram Burgard
Department of Computer Science, University of Freiburg
79110 Freiburg, Germany

Abstract

Mobile robots that provide service to people can carry out their tasks more efficiently if they know where the people are. In this paper we present an approach to actively maintain a probabilistic belief about the current locations of people in the environment of a mobile robot. We assume that the robot is equipped with knowledge about typical motion behaviors of people in form of Hidden Markov Models (HMMs), which are updated based on vision and laser information. While the robot is carrying out its task it applies a decision-theoretic approach to actively select points in the environment that are expected to provide information about the positions of people. Experimental results obtained with a mobile robot in a typical office environment illustrate that our method decreases the uncertainty about the positions of people compared to passive approaches which do not consider additional observation actions.

1 Introduction

Service robots are envisioned to coexist with humans and to fulfill various tasks such as transportation [11], cleaning [5], entertainment [2], and assistance of people in their everyday activities [1, 15]. In many of such tasks it is useful to know the current locations of the people in the environment. For example, a robot that knows which rooms are currently empty can carry out its cleaning tasks without disturbing anyone. Furthermore, a robot that anticipates where a person might be going can generate motion actions that avoid interferences with the person. Finally, knowledge about where people currently are allows the robot to more efficiently carry out delivery tasks since the number of detours is reduced.

In this paper we investigate the problem of how a robot can effectively maintain a probabilistic belief about the positions of the people in the environment. We assume that the robot possesses information about typical motion behaviors of people in form of Hidden Markov Models (HMMs) [3]. These HMMs are updated based on sensory input while the robot is carrying out its tasks. We use a decision-theoretic approach to identify whether the robot should add observation actions to its task. Throughout this paper an observation action corresponds to moving to a place in the environment and obtaining a sensor measurement there.

The problem of localizing people in the environment of mobile robots was studied intensively in the past. For example, several authors concentrated on the question of how to track multiple moving

targets in the vicinity of the robot [24, 8, 16, 12]. Nguyen et al. [19] proposed to use an Abstract Hidden Markov mEmory Model (AHMEM) to maintain a belief about the positions of people and to infer their intentions. All those authors mainly focused on the question of how to represent and update the belief about the individual targets being tracked and did not consider actions of the robot so as to maximize the tracking performance.

González-Baños et al. [13, 7] considered the problem of maximizing the visibility of a moving target which is followed by a robotic observer. La Valle et al. [14] studied the problem of localizing an unknown object in a workspace. Additionally, several multi-robot systems were developed that keep track of moving objects [18, 9, 20] or surround a moving target [21]. In the work presented by Rosencrantz et al. [23] a team of robots tries to locate and tag “enemies” which are not always in their perceptual field of view. They used variable-dimension particle filters to track the location of the moving objects. The movements for the observer robots are coordinated so that the information gain is maximized and the search time minimized.

All these approaches are either passive in the sense that they just maintain a belief about the positions of the targets being tracked or are reactive and generate short-term plans to maintain or maximize the visibility of the objects being tracked. In this paper, in contrast, we consider the problem of actively maintaining an accurate belief about the positions of people while the robot has to carry out navigation tasks such as office delivery. We use learned Hidden Markov Models [3] representing typical motion behaviors to maintain beliefs about the locations of the people and update these models upon sensory input. To detect and identify people our approach combines laser range data with vision information. According to the nature of HMMs, our system is able to exploit even negative information to update its belief about potential positions of people. A decision-theoretic approach is used to decide whether the robot should integrate observation actions into its plans. This way, our robot can effectively minimize the uncertainty about the positions of people while it is carrying out its tasks.

The paper is organized as follows. In the next section we explain the structure of the HMMs that we use to estimate the locations of people. In Section 3 we describe how our robot uses its sensor information to identify people in the environment. Section 4 introduces our strategy how to decide if and which observation action should be executed to update the robot’s belief. In Section 5 we present several experiments illustrating the robustness of our approach to maintain an accurate belief about the positions of people using laser and vision data with a mobile robot.

2 Hidden Markov Models for Typical Motion Behaviors of People

Our approach assumes that the typical motion behaviors of the people are given in the form of Hidden Markov Models [22]. People usually do not permanently move. Rather they typically move between so-called *resting places*. These are places where the people frequently stay for a while. Therefore the HMM we use distinguishes two types of states [3]. The first class of states are the resting places. The second class are *intermediate states* which lie on typical trajectories that the people follow when walking from resting place to resting place. The transition probabilities of the intermediate states model average walking speed, whereas the transition probabilities for the resting places are computed based on a statistics about the average time period that elapses until the person starts to move on a particular trajectory after arriving at the corresponding resting place.

Figure 1 depicts the structure of the HMM for our office environment. The numbered squares indicate the eight resting places and the small circles on the trajectories are the intermediate states. The whole HMM consists of 966 states.




Figure 1. Possible transitions of the Hidden Markov Model derived from learned motion patterns in our laboratory. The numbered squares indicate the eight resting places and the small circles on the trajectories are the intermediate states.

3 Person Detection and Identification

To keep track of multiple persons in an environment, one in principle would have to maintain a belief about the joint state space of all persons. This approach, however, is usually not feasible since the complexity of the state estimation problem grows exponentially with the number of persons or dimensions of the state space. Therefore we approximate the posterior by factorizing the belief about the joint state space and consider independent beliefs about the states of all persons. Each such belief is represented by a separate HMM. To maintain the individual beliefs we need to be able to update the HMMs for the persons based on observations made by the robot, which requires the ability to reliably detect people and to identify them. To achieve this, our current system combines laser and vision information [3]. To detect people in the laser-range scans obtained with the robot our system extracts features which are local minima that correspond to the people's legs. We also need to be able to identify a person in order to appropriately update the belief about the location of that person. To achieve this we employ the vision system of our robot and learn an image database beforehand. For each person this database contains one histogram which is built from 20 images. To identify a person, we proceed as follows. Every time the laser-based people detection system reports a feature in the field of view of the camera, an image is collected and the following three steps are applied. *Segmentation*: We determine a rectangular area in the camera image which contains the person. To determine this area we use a perspective projection to map the 3D position of the person in world coordinates to 2D image coordinates. *Feature extraction*: We compute a color histogram for the area selected in the previous step. *Database matching*: To determine the likelihood of a particular person, we compare the histogram computed in step 2 to all prototypes existing in the database.

Thus, each observation z_j^t consists of the position y_j^t of a feature provided by a laser-based people detector and a similarity measure $H(q_j^t, \pi_i)$ between the query color histogram q_j^t of the corresponding segment in the camera image and the database histogram of person i , which has been learned before. To compare a given query histogram q_j^t with a prototype π_i in the database we use the normalized intersection norm [26].

Let ξ_i^t denote the state of person i at time t . Whenever a new observation is obtained we apply Joint Probabilistic Data Association Filters [4] and integrate the individual features according to the assignment probability λ_{ji} that feature j corresponds to person i :

$$P(\xi_i^t | z^{(1:t)}) = \eta \sum_{j=0}^{S^t} \lambda_{ji} P(z_j^t | \xi_i^t) P(\xi_i^t | z^{(1:t-1)}). \quad (1)$$

Here η is a normalization factor, S^t is the number of features detected at time t , and $z^{(1:t)}$ denotes the sequence of all measurements up to time t (see also Schulz et al. [24]).

We compute the likelihood $P(z_j^t | \xi_i^t)$ according to:

$$P(z_j^t | \xi_i^t) = P(y_j^t, H(q_j^t, \pi_i) | \xi_i^t) = H(q_j^t, \pi_i) P(y_j^t | \xi_i^t). \quad (2)$$

Here $P(y_j^t | \xi_i^t)$ is the probability that the laser-based people detection system reports a feature detection at location y_j^t given that the person is in state ξ_i^t . We determine this quantity using a mixture of a uniform distribution and a bounded Gaussian with mean y_j^t . Note that we also take into account visibility constraints, i.e., states that are occluded are regarded as states outside the bounded Gaussian. In the case that no feature has been obtained for a person we use the likelihood of false negative observations for such states that are within the range of the robot's sensors. For all other states we use the average likelihood that the robot does not detect a person given it is outside the sensor range.

4 Computing the Expected Utility of Observation Actions

Since the robot becomes uncertain about the position of a person when it has not been observing the person for a longer period of time, it should consider to actively perform observation actions to update its belief. In this context two aspects are relevant. On one hand, the information gain should be as large as possible, and on the other hand, the cost of performing observation actions should be minimized.

To determine the uncertainty in the belief about the positions of the people we consider the entropy of the posteriors. The entropy H of the posterior Bel_j about possible states ν_1, \dots, ν_{N_j} of person j is a general measure for the uncertainty and is defined as:

$$H(Bel_j) = - \sum_{i=1}^{N_j} Bel_j(\nu_i) \cdot \log Bel_j(\nu_i). \quad (3)$$

H is maximal in case of a uniform distribution. The minimal value zero is obtained if the robot is absolutely certain about the current position of person j .

To take into account the information provided by the sensors of the robot, we compute the expected information gain which is the expected change of entropy given that the robot obtains sensor information. The information gain for the posterior Bel_j given an observation z is defined as:

$$I(Bel_j | z) = H(Bel_j) - H(Bel_j | z). \quad (4)$$

Here $H(Bel_j | z)$ is the entropy of the posterior about the position of person j after integrating the observation z .

Note that the problem considered here – choosing the optimal action sequence – can be regarded as a partially observable Markov decision process (POMDP) problem (see [10] for a comprehensive overview). Since solving the POMDP for applications of the size considered here is not feasible in practice we follow an approach that makes several simplifying assumptions and has been applied successfully for a similar problem in the past [6]. First, we use a restricted set of potential observation actions that the robot can carry out. Since we assume that the people stay at the resting places most of the time we only consider observation actions at viewpoints for the resting places. Furthermore, we consider only one observation per task and do not consider all potential measurements perceived by the robot while it is carrying out its task. A task is defined as the action of driving to the next target location of the robot.



Figure 2. The left image shows the visibility area for resting place 6. The right image shows for two scenarios the detours to the viewpoints to which the robot moves. The solid lines indicate the direct paths to the target locations of the robot and the dashed lines are the detours to the viewpoints.

Additionally, we take the possibility into account that the robot can observe a resting place whenever it arrives at the final location of its current task. The case that no part of the HMM can be observed after the robot arrived at its goal can be regarded as a special case of this.

Since we do not know what the robot will perceive when it has executed its task a , we have to sum over all possible observations z_a to compute the expected information gain for Bel_j :

$$E(Bel_j|a) = \sum_{z_a} p(z_a | Bel_j) \cdot I(Bel_j | z_a). \quad (5)$$

To efficiently compute the likelihood $p(z_a | Bel_j)$ we do not integrate over all possible measurements. Instead, we consider abstract observations namely that the robot identifies/does not identify the person given it is at the observed resting place or not. The corresponding likelihoods are the average detection respectively failure rates. The overall expected information gain $E(a)$ for the task a is given by the sum of the individual expected information gains for the posteriors of all J persons after executing a :

$$E(a) = \sum_{j=1}^J E(Bel_j|a). \quad (6)$$

The expected utility of a can now be defined as:

$$EU(a) = r(a) - c(a) + E(a). \quad (7)$$

Here, r specifies a reward function which depends on the utility of finishing a , whereas $c(a)$ are the cost of executing a .

During the execution of its current task, we allow the robot to consider one additional observation of a resting place. To reduce the complexity we compute viewpoints for the resting places. From these viewpoints the robot can observe if a person is currently staying at the corresponding resting place.

To compute the viewpoints we proceed as follows: We perform two deterministic value iterations [25] in the static 2D occupancy grid map [17] of the environment¹: one with the target position of the current task as starting location and one with the robot's current position as starting location.

¹Each cell $\langle x, y \rangle$ of the occupancy grid map stores the probability $p_{occ}(\langle x, y \rangle)$ that the corresponding area in the environment is occupied. The cost for traversing a cell $\langle x, y \rangle$ is proportional to its occupancy probability $p_{occ}(\langle x, y \rangle)$.

For each resting place l we compute a visibility area, i.e. the set of grid cells from which l is visible. The left image of Figure 2 shows for example the visibility area of resting place 6. The viewpoint for a resting place l is then defined as the cell in its visibility area which has the lowest move costs. To determine the move cost for a cell $\langle x, y \rangle$ we can simply add the costs for reaching $\langle x, y \rangle$ from the start location and for reaching the final location from $\langle x, y \rangle$. Additionally we add a penalty term corresponding to the cost imposed by rotating the robot towards l . The right image of Figure 2 depicts for two scenarios the detours of the robot to viewpoints of resting place 6.

The action of moving to the viewpoint corresponding to a resting place l to get an observation is referred to as a_1 and the action of driving from l to the original target location is referred to as a_2 . Let $a_l = a_1 \oplus a_2$ be the whole action consisting of a_2 executed after a_1 . The expected utility of a_l is given by:

$$EU(a_l) = r(a_1 \oplus a_2) - c(a_1 \oplus a_2) + E(a_1 \oplus a_2). \quad (8)$$

Note that both a_1 and a_2 depend on the resting place l . However, to enhance readability we omitted the argument l . The expected information gain for $a_1 \oplus a_2$ can be computed as:

$$E(a_1 \oplus a_2) = \sum_{j=1}^J \sum_{z_{a_1}} \sum_{z_{a_2}} p(z_{a_1} | Bel_j) \cdot p(z_{a_2} | Bel_j) \cdot I(Bel_j | z_{a_1}, z_{a_2}). \quad (9)$$

During the execution of its current task a the robot moves to the viewpoint corresponding to resting place l^* with

$$l^* = \operatorname{argmax}_{l \in \{l_1, \dots, l_N\}} EU(a_l) \quad (10)$$

whenever $EU(a_{l^*}) > EU(a)$. Here l_1, \dots, l_N are the resting places.

5 Experimental Results

The approach described above has been implemented and tested using a B21r robot equipped with a laser range scanner and a stereo camera system (see left image in Figure 5). The states and possible transitions of the HMM we used to track the positions of the persons are depicted in Figure 1. The update of all individual HMMs and the computation of the viewpoint with the maximum expected utility is done in 0.5 seconds on a 3 GHz Pentium 4. Thus, the robot considered additional observation actions at a frequency of 2 Hz. We assume that the reward $r(a)$ is equal for all actions. The results obtained illustrate that our system allows the robot to actively perform observation actions and to use these actions to reduce its uncertainty about the current positions of the people in its environment.

5.1 Performing an Observation Action During Task Execution

The goal of the first experiment is to illustrate that our algorithm can effectively guide the robot to viewpoints that provide information about positions of people when needed. The task of the robot was to move from the position marked with $t = 0$ in Figure 3 to resting place 4. In the initial situation the viewpoint of resting place 3 had the highest expected utility because the robot was uncertain about the current position of person X (see initial belief in the left image of Figure 4) and because the additional move costs for viewpoint 3 were very low. Therefore the robot decided to observe resting place 3 at time step $t = 50$. The robot detected person X and as can be seen in the left image of




Figure 3. The robot decides to stop at the viewpoint to check whether person X is in its room (resting place 3) or not while it is executing a task.




Figure 4. The left image shows the evolution of the probability of person X to be at different resting places over time for the experiment shown in Figure 3. As can be seen at time step $t = 50$ when the robot detects person X at resting place 3 it updates its belief accordingly. The right image shows the corresponding evolution of the probability when the robot does not stop at the viewpoint to perform an observation action.

Figure 4 updated its belief accordingly. If, in contrast to this, the robot in the same situation does not perform an observation action at viewpoint 3 and moves directly to its target location, its belief about the position of person X would not improve over time. This fact is illustrated in the right image of Figure 4. In all the figures we only show posteriors of the relevant resting places to enhance readability. In a similar experiment, in which the robot was quite certain that person X was in its office, it did not stop at the viewpoint because the expected utility was not high enough.

5.2 Actively Searching for a Person

The second experiment has been designed to illustrate that the robot can deal with ambiguities and that it can increase its certainty about the positions of the person by integrating negative information. Here the robot was standing in the middle of the corridor looking to the east and currently had no task to execute. At around time step $t = 20$ the robot observed a person walking to the east. A scene overview is depicted in the left image of Figure 5. According to the camera information the detected person was most likely person B , who previously had been staying at resting place 4. The resulting posterior about the position of person B after integrating the observation sequence is depicted in the left image of Figure 6. The size of the squares of the intermediate states of the HMM represents the probability that the person is currently in the corresponding state. Similarly the resting places




Figure 5. The left image shows person *B* walking down the corridor. It enters the room containing resting place 7 and walks to resting place 6 in the neighbor room. Since the robot is uncertain where person *B* is going to after he moved out of the field of view, the robot decides to search for it. The viewpoint corresponding to resting place 7 has the highest expected utility and therefore the robot moves there to perform an observation action (center image). The right image shows the evolution of the probability of person *B* to be at different resting places over time. When the robot does not detect a person at resting place 7 (time step $t = 73$) it infers that person *B* is probably staying at resting place 6.

are labeled with the probability that the person is staying at this particular place. The circle labeled person *B* corresponds to the position of the person provided by the laser-based people detection system. At around time step $t = 30$ person *B* disappeared out of the field of view of the robot and walked through the office containing resting place 7 to resting place 6. Since the robot was uncertain to which resting place person *B* was going to, the robot decided to search for it.

According to the transition probabilities of the resting places, the robot believed that the person most likely walked to resting place 6. Still, the probability of resting place 7 was quite high (see right image of Figure 5). Since the viewpoint corresponding to resting place 7 had the highest expected utility, the robot decided to move there to perform the corresponding observation action (see center image of Figure 5). In this example, the robot did not observe person *B* at resting place 7 so that, after the update of the HMM, the person was most likely at resting place 6. The right image of Figure 5 shows the evolution of the belief about the position of the person during this experiment. As can be seen, the probability of person *B* to be at resting place 6 rapidly increased at time step $t = 73$ when the robot checked resting place 7 and did not detect it there.

5.3 Keeping Track of Multiple Persons

The final experiment described in this section is designed to illustrate that the robot can actively maintain a belief about the positions of multiple persons. In this particular experiment the robot was keeping track of two persons and was standing at the same place as in the beginning of the previous experiment.

Initially the robot believed that persons *A* and *B* were most likely at resting place 4. At around time step $t = 15$ the robot observed one person walking to the east along the corridor and entering the office containing resting place 7. Since the similarity measures between the extracted segment in the camera image and the database histograms of person *A* and *B* were ambiguous, the robot was rather uncertain which person it had observed. The similarity between the corresponding segment in the camera image and the database image of person *B* was only slightly higher than the similarity to the database image of person *A*. More precisely, the likelihood that the detected person was person *B* was 0.57 and the likelihood of person *A* was 0.43. Therefore, the probability of person *B*, who actually

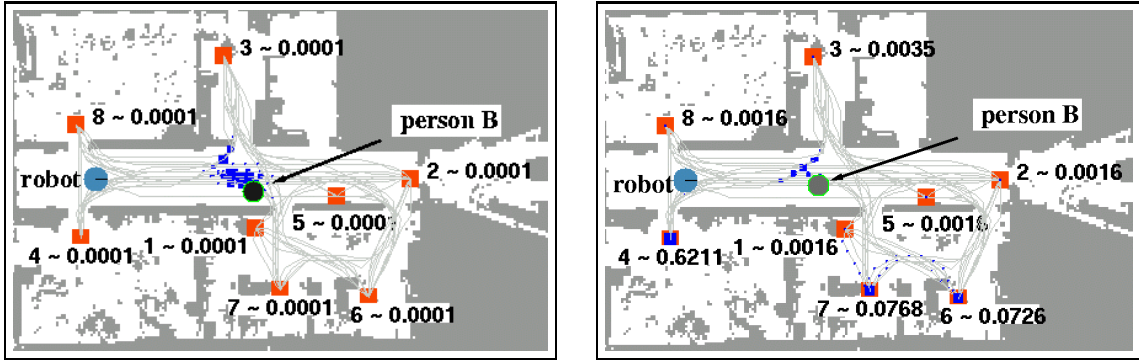


Figure 6. The robot detects that person B walks away through the corridor. Shown on the left is the belief about the position of person B after integrating part of the observation sequence. The size of the squares of the intermediate states of the HMM represents the probability that the person is currently in the corresponding state. In the experiment shown on the right image the robot detects a person in the corridor. Since the robot is rather uncertain whether the person is A or B the probability that person B still stays at resting place 4 remains high.

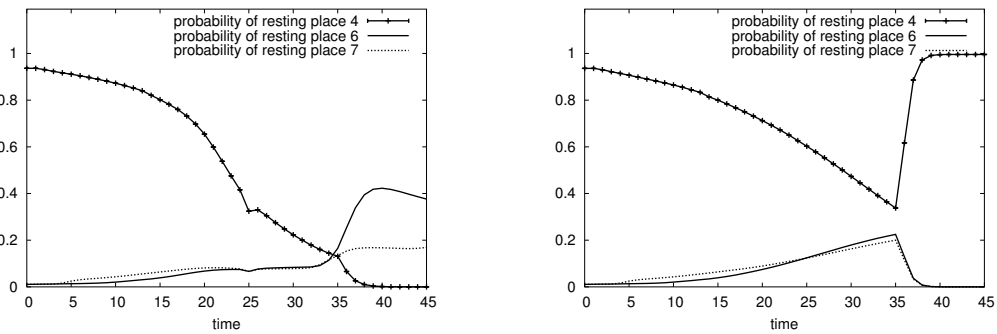


Figure 7. The left image shows the evolution of the probability of person B to be at different resting places over time for the case that the robot is uncertain whether the person it detected was person A or person B and therefore decides to inspect resting place 4. When the robot observes person A there the probabilities of resting places 6 and 7 immediately increase for person B (time step $t = 35$). The right image shows that at the same time the probability of person A to be at resting place 4 seriously increases.

walked down the corridor, to be at resting place 4 remained high. The right image of Figure 6 depicts the posterior about the position of person B after integrating part of the observation sequence. At around time step $t = 25$ the robot decided to turn to resting place 4 to check which person was still there. The robot identified person A at time step $t = 35$ and updated its belief accordingly. As can be seen in the left image of Figure 7 the probabilities that person B was at the resting places 6 and 7 immediately increased after the inspection of resting place 4. The right image shows the evolution of the belief about the position of person A . As can be seen from the graph the robot was now very certain that person A was at resting place 4.

6 Conclusions and Future Work

In this paper we considered the problem of actively maintaining an accurate belief about the positions of multiple persons in the environment of a mobile robot. Our approach represents the probabilistic beliefs about the positions of people by Hidden Markov Models and updates these models using Joint Probabilistic Data Association Filters. It uses a decision-theoretic approach to determine observation actions that are carried out while the robot is executing its tasks. The utility of an observation action is computed by trading-off move costs and the expected reduction of uncertainty.

Our approach has been implemented on a real robot and evaluated using data provided by a laser-range sensor and cameras. Experimental results demonstrate that our algorithm generates effective actions that seriously reduce the uncertainty in the belief about the positions of people.

In the future, we plan to investigate whether our models can be applied to predict the behaviors of people in other application scenarios like for example shopping malls or museums. In these environments the people also show typical behavior and frequently stop at certain places. The difficulty here is that the environments are more crowded. Since in such kind of environments the people are often moving in groups it will probably be helpful to track them as one group and not as individuals.

Acknowledgments

This work has partly been supported by the German Science Foundation (DFG) under contract number SFB/TR-8. We would like to thank Grzegorz Cielniak for providing us with the vision-based people identification system and Cyrill Stachniss for helping us to carry out experiments with our mobile robot Albert.

References

- [1] H. Asoh, S. Hayamizu, I. Hara, Y. Motomura, S. Akaho, and T. Matsui. Socially embedded learning of office-conversant robot Jijo-2. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1997.
- [2] W. Burgard, A. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1-2), 1999.
- [3] G. Cielniak, M. Bennewitz, and Burgard. Where is ...? Learning and utilizing motion patterns of persons with mobile robots. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2003.
- [4] I. Cox. A review of statistical data association techniques for motion correspondence. *International Journal of Computer Vision*, 10(1):53–66, 1993.
- [5] H. Endres, W. Feiten, and G. Lawitzky. Field test of a navigation system: Autonomous cleaning in supermarkets. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 1998.

- [6] D. Fox, W. Burgard, and S. Thrun. Active markov localization for mobile robots. *Robotics and Autonomous Systems*, 25:195–207, 1998.
- [7] H. González-Baños, C.-Y. Lee, and J.-C. Latombe. Real-time combinatorial tracking of a target moving unpredictably among obstacles. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2002.
- [8] B. Jensen, R. Philippsen, and R. Siegwart. Narrative situation assessment for human-robot interaction. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2003.
- [9] B. Jung and G. Sukhatme. A region-based approach for cooperative multi-target tracking in a structural environment. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2002.
- [10] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101, 1998.
- [11] S. King and C. Weiman. Helpmate autonomous mobile robot navigation system. In *Proc. of the SPIE Conference on Mobile Robots*, pages 190–198, Boston, MA, 1990.
- [12] B. Kluge, C. Köhler, and E. Prassler. Fast and robust tracking of multiple moving objects with a laser range finder. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2001.
- [13] S. Lavalle, H. González-Banos, G. Becker, and J.-C. Latombe. Motion strategies for maintaining visibility of a moving target. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 1997.
- [14] S. LaValle, D. Lin, L. Guibas, J. Latombe, and R. Motwani. Finding an unpredictable target in a workspace with obstacles. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 1997.
- [15] M. Montemerlo, J. Pineau, N. Roy, S. Thrun, and V. Verma. Experiences with a mobile robotic guide for the elderly. In *Proc. of the AAAI National Conference on Artificial Intelligence*, 2002.
- [16] M. Montemerlo, S. Thrun, and W. Whittaker. Conditional particle filters for simultaneous mobile robot localization and people-tracking. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2002.
- [17] H. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 116–121, 1985.
- [18] R. Murrieta-Cid, H. González-Baños, and B. Tovar. A reactive motion planner to maintain visibility of unpredictable targets. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2002.
- [19] N. Nguyen, H. Bui, S. Venkatesh, and G. West. Recognising and monitoring high-level behaviours in complex spatial environments. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003.
- [20] L. Parker. Cooperative motion control for multi-target observation. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1997.
- [21] P. Pirjanian and M. Matarić. Multi-robot target acquisition using multiple objective behaviour coordination. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2000.
- [22] L. Rabiner and B. Juang. An introduction to hidden markov models. *IEEE ASSP Magazine*, 3(1):4–16, 1986.
- [23] M. Rosencrantz, G. Gordon, and S. Thrun. Locating moving entities in dynamic indoor environments with teams of mobile robots. In *Proc. of the Second Joint International Conference on Autonomous Agents & Multi Agent Systems (AAMAS)*, 2003.
- [24] D. Schulz, W. Burgard, D. Fox, and A. Cremers. People tracking with a mobile robot using sample-based joint probabilistic data association filters. *International Journal of Robotics Research (IJRR)*, 22(2):99–116, 2003.
- [25] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.
- [26] M. Swain and D. Ballard. Color indexing. *International Journal of Computer Vision*, 7(1), 1991.