

Learning Motion Patterns of Persons for Mobile Service Robots

Maren Bennewitz[†] Wolfram Burgard[†] Sebastian Thrun[‡]

[†]Department of Computer Science, University of Freiburg, 79110 Freiburg, Germany

[‡]School of Computer Science, Carnegie Mellon University, Pittsburgh PA, USA

Abstract

We propose a method for learning models of people’s motion behaviors in an indoor environment. As people move through their environments, they do not move randomly. Instead, they often engage in typical motion patterns, related to specific locations that they might be interested in approaching and specific trajectories that they might follow in doing so. Knowledge about such patterns may enable a mobile robot to develop improved people following and obstacle avoidance skills. This paper proposes an algorithm that learns collections of typical trajectories that characterize a person’s motion patterns. Data, recorded by mobile robots equipped with laser range finders, is clustered into different types of motion using the popular expectation maximization algorithm, while simultaneously learning multiple motion patterns. Experimental results, obtained using data collected in a domestic residence and in an office building, illustrate that highly predictive models of human motion patterns can be learned.

1 Introduction

Whenever mobile robots are designed to operate in populated environments, they need to be able to perceive the people in their environment and to adapt their behavior according to the activities of the people. The knowledge of typical behaviors can be used in several ways to improve the behavior of the system. For example, it allows a robot to adapt its velocity to the speed of people in the environment and it enables a robot to choose trajectories that minimize the risk of collisions with people. In this paper we consider a specific problem in the context of a nursing robot project [13]. The goal of this project is the development of intelligent service robots that can assist people in their daily living activities. One aspect in this context is to learn typical behaviors of the persons in order to know, where the person currently is or where it is currently going to.

Recently, a variety of service robots were developed that are designed to operate in populated environments. These robots, for example, are deployed in hospitals [7], museums [3], office buildings [1], and department stores [4], where they perform various services, e.g., deliver, educate, entertain [15] or assist people [14, 9]. Additionally, a variety of techniques has been developed that allows a robot to estimate the positions of people in its vicinity or to adapt its behavior accordingly. For ex-

ample, the techniques described in [16] are designed to track multiple persons in the vicinity of a robot. The approach presented in [17] uses a given probabilistic model of typical motion behaviors in order to predict future poses of the persons. [6] present an approach to improve the behavior of a robot by following the activities of a teacher. The system described in [8] uses a camera to estimate where persons typically walk and adapts the trajectory of the robot appropriately. [18] apply a Hidden-Markov-Model to predict the motions of moving obstacles in the environment of a robot. [10] present a system that is able to keep track of a moving target even in the case of possible occlusions by other obstacles in the environment. All the techniques described above assume the existence of a model of the motion behaviors. Our approach, in contrast, is able to learn such models and to use the learned models for the prediction of the peoples movements. The technique described in [2] uses an Abstract Hidden-Markov-Model to learn and to predict motions of a person. This approach assumes that all trajectories are already clustered into the corresponding motion behaviors during the learning phase. Our method extends this approach as it determines both, the clustering and the corresponding motion behaviors.

In this paper we present an approach that allows a mobile robot to learn probabilistic motion patterns of persons. We use the popular EM-algorithm to simultaneously cluster trajectories belonging to one motion behavior and to learn the characteristic motions of this behavior. We apply our technique to data recorded with mobile robots that are equipped with laser-range finders. Furthermore, we demonstrate how the learned models can be used to predict the trajectory of a person in the natural environment.

This paper is organized as follows. In the next section, we present the probabilistic representation of the motion patterns and describe how to learn them using the expectation maximization algorithm. In Section 3 we describe our application based on data recorded with mobile robots that are equipped with laser-range finders. Section 4 presents experimental results regarding the learning process as well as regarding the prediction accuracy of the learned models.

2 Learning Motion Patterns

Our approach to discovering typical motion patterns of people is strictly statistical, using the popular EM-algorithm to

find different types of activities that involve physical motion throughout the natural environment. The input to our routine is a collection of trajectories $d = \{d_1, \dots, d_N\}$ (called: the data). The output is a number of different types of motion patterns $\theta = \{\theta_1, \dots, \theta_M\}$ a person might exhibit in their natural environment. Each trajectory d_i consists of a sequence

$$d_i = \{x_i^1, x_i^2, \dots, x_i^T\} \quad (1)$$

of positions x_i^t covered by the person. In our current system, these positions are computed based on a grid-based discretization of the environment: Each x_i^t represents the position of the cell covered by the person after t steps. Accordingly, x_i^1 is the first cell covered by the person and x_i^T is the final destination. Throughout this paper we assume that all trajectories d_i have the same length. In our current system we choose T as the maximum length of all trajectories. Trajectories of length $T' < T$ are extended to length T by adding the final location of that trajectory for exactly $T - T'$ times.

2.1 Motion Patterns

We begin with the description of our model of motion patterns, which is subsequently estimated from data using EM. Within this paper we assume that a person engages in M different types of motion patterns. A motion pattern, denoted θ_m with $1 \leq m \leq M$, is represented by probability distributions $p(x | \theta_m^t)$ specifying the probability that the person is at location x after t steps given that he or she is engaged in this motion pattern. Accordingly, we calculate the likelihood of a trajectory d_i under the m -th motion model θ_m as

$$p(d_i | \theta_m) = \prod_{t=1}^T p(x_i^t | \theta_m^t). \quad (2)$$

2.2 Expectation Maximization

In essence, our approach seeks to identify a model θ that maximizes the likelihood of the data. To define the likelihood of the data under the model θ , it will be useful to introduce a set of *correspondence variables*, denoted c_{im} . Here i is the index of a trajectory d_i , and m is the index of a motion model θ_m . Each correspondence c_{im} is a binary variable, that is, it is either 0 or 1. It is 1 if and only if the i -th trajectory corresponds to the m -th motion pattern. If we think of the motion model as a specific motion activity a person might engage in, c_{im} is 1 if person was engaged in motion activity m in trajectory i .

In the sequel, we will denote the set of all correspondence variables for the i -th data item by c_i , that is, $c_i = \{c_{i1}, \dots, c_{iM}\}$. For any data item i , the fact that exactly one correspondence is 1 translates into the following:

$$\sum_{m=1}^M c_{im} = 1. \quad (3)$$

Throughout this paper we assume that each motion pattern is represented by T Gaussian distributions with a fixed standard

deviation σ . Accordingly, the application of EM leads to an extension of the k-Means Algorithm (see e.g. [12]) to trajectories. Given the individual Gaussians for a model θ we can compute the joint likelihood of a single trajectory d_i and its correspondence vector c_i as follows:

$$p(d_i, c_i | \theta) = \prod_{t=1}^T \prod_{m=1}^M \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2\sigma^2} c_{im} \|x_i^t - \mu_m^t\|^2}. \quad (4)$$

Thereby, we exploit the fact that only one of the correspondence variables c_{im} in the inner product is 1, and all others are 0. Accordingly, the total likelihood over all values of i is given by the product of the individual joint probabilities:

$$p(d, c | \theta) = \prod_{i=1}^N \left(\prod_{t=1}^T \prod_{m=1}^M \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2\sigma^2} c_{im} \|x_i^t - \mu_m^t\|^2} \right). \quad (5)$$

Since the logarithm is a monotonic function we can maximize the log likelihood instead of the likelihood. The logarithm of (5) is given by:

$$\ln p(d, c | \theta) = \sum_{i=1}^N \left(T \cdot M \cdot \ln \frac{1}{\sqrt{2\pi}\sigma} - \frac{1}{2\sigma^2} \cdot \sum_{t=1}^T \sum_{m=1}^M c_{im} \|x_i^t - \mu_m^t\|^2 \right). \quad (6)$$

Finally, we notice that we are not really interested in the log likelihood of the correspondence variables c , since those are not observable in the first place. The common approach is to integrate over them, that is, to optimize the expected log likelihood $E_c[\ln p(d, c | \theta) | \theta, d]$ instead which, according to (6), is

$$\begin{aligned} E_c[\ln p(d, c | \theta) | \theta, d] &= E_c \left[\sum_{i=1}^N \left(T \cdot M \cdot \ln \frac{1}{\sqrt{2\pi}\sigma} - \frac{1}{2\sigma^2} \cdot \sum_{t=1}^T \sum_{m=1}^M c_{im} \|x_i^t - \mu_m^t\|^2 \right) \middle| \theta, d \right]. \quad (7) \end{aligned}$$

Since the expectation is a linear operator we can move it inside the expression, so that we finally get:

$$\begin{aligned} E_c[\ln p(d, c | \theta) | \theta, d] &= \sum_{i=1}^N \left(T \cdot M \cdot \ln \frac{1}{\sqrt{2\pi}\sigma} - \frac{1}{2\sigma^2} \cdot \sum_{t=1}^T \sum_{m=1}^M E[c_{im} | \theta, d] \|x_i^t - \mu_m^t\|^2 \right), \quad (8) \end{aligned}$$

where $E[c_{im} | \theta, d]$ depends on the model θ and the data d .

Unfortunately, optimizing (8) is not an easy endeavor. EM is an algorithm that iteratively maximizes expected log likelihood functions by optimizing a sequence of lower bounds.

In particular, it generates a sequence of models, denoted $\theta^{[1]}, \theta^{[2]}, \dots$ of increasing log likelihood.

Mathematically, the standard method is to turn (8) in a so-called Q -function which depends on two models, θ and θ' :

$$Q(\theta' | \theta) = E_c[\ln p(d, c | \theta') | \theta, d]. \quad (9)$$

In accordance with (8), this Q -function is factored as follows:

$$\begin{aligned} Q(\theta' | \theta) &= \sum_{i=1}^N \left(T \cdot M \cdot \ln \frac{1}{\sqrt{2\pi\sigma}} \right. \\ &\quad \left. - \frac{1}{2\sigma^2} \cdot \sum_{t=1}^T \sum_{m=1}^M E[c_{im} | \theta, d] \|x_i^t - \mu_m^t\|^2 \right). \quad (10) \end{aligned}$$

The sequence of models is then given by calculating

$$\theta^{[j+1]} = \underset{\theta'}{\operatorname{argmax}} Q(\theta' | \theta^{[j]}) \quad (11)$$

starting with some initial model $\theta^{[0]}$. Whenever the Q -function is continuous as in our case, the EM algorithm converges at least to a local maximum.

In particular, the optimization involves two steps: calculating the expectations $E[c_{im} | \theta^{[j]}, d]$ given the current model $\theta^{[j]}$, and finding the new model $\theta^{[j+1]}$ that has the maximum expected likelihood under these expectations. The first of these two steps is typically referred to as the E-step (short for: expectation step), and the latter as the M-step (short for: maximization step).

To calculate the expectations $E[c_{im} | \theta^{[j]}, d]$ we apply Bayes rule, obeying independence assumptions between different data trajectories:

$$\begin{aligned} E[c_{im} | \theta^{[j]}, d] &= p(c_{im} | \theta^{[j]}, d) \\ &= p(c_{im} | \theta^{[j]}, d_i) \\ &= \eta p(d_i | c_{im}, \theta^{[j]}) p(c_{im} | \theta^{[j]}) \\ &= \eta' p(d_i | \theta_m^{[j]}), \quad (12) \end{aligned}$$

where the normalization constants η and η' ensure that the expectations sum up to 1 over all m . If we combine (2) and (12) exploiting the fact that the distributions are represented by Gaussians we obtain:

$$E[c_{im} | \theta^{[j]}, d_i] = \eta' \prod_{t=1}^T e^{-\frac{1}{2\sigma^2} \|x_i^t - \mu_m^t\|^2}. \quad (13)$$

Finally, the M-step calculates a new model $\theta^{[j+1]}$ by maximizing the expected likelihood. Technically, this is done by computing for every motion pattern m and for each time step t a new mean $\mu_m^{t[j+1]}$ of the Gaussian distribution. Thereby we consider the expectations $E[c_{im} | \theta^{[j]}, d]$ computed in the E-step:

$$\mu_m^{t[j+1]} = \frac{\sum_{i=1}^N E[c_{im} | \theta^{[j]}, d] x_i^t}{\sum_{i=1}^N E[c_{im} | \theta^{[j]}, d]} \quad (14)$$

2.3 Monitoring Convergence and Local Maxima

The EM-algorithm is well-known to be sensitive to local maxima in the search [5, 11]. In the context of identifying motion patterns, a typical local maximum involves situations in which different types of trajectories are, with high probability, associated with the same model component θ_m . In such cases, the motion patterns may never develop into a clear model of people's motion, and specific trajectories may never be explained well with any of the model components.

Luckily, such cases can be identified during the optimization. Our approach continuously monitors two types of occurrences:

1. Low data likelihood: If a trajectory d_i has low likelihood under the model θ , this is an indication that no appropriate motion pattern has yet been identified that represents this trajectory.
2. Low motion pattern utility: Our second criterion involves testing the utility of a motion patterns. The aim of this criterion is to discover multiple model component that basically represent the same people motion. To detect such cases, the total data log likelihood is calculated with and without a specific model component θ_m . Technically, this involves executing the E step twice, once with and once without θ_m . If the difference in the overall data likelihood is smaller than a pre-specified threshold, the effect of removing θ_m from the model is negligible. This indicates a case where a similar motion pattern exists and the one at hand is a duplicate.

Whenever the EM appears to have converged, our approach extracts those two statistics and considers "resetting" individual model components. In particular, if a low data likelihood trajectory is found, a new model component is introduced that is initialized using this very trajectory. Conversely, if a model component with low utility is found, it is eliminated from the model.

In our experiments, we found this selective restarting and elimination strategy extremely effective in escaping local maxima. As indicated in the experimental results section below, all our experiments converged to a model that clustered trajectories into categories 100% identical to those prescribed by us manually. Without this mechanism, the EM frequently got stuck in local maxima and generated models that were significantly less predictive of human behavior.

3 Laser-based Implementation

The EM-based learning procedure has been implemented for data acquired with laser-range finders. To acquire the data we used three Pioneer I robots (see left image of Figure 1) which we installed in the environments. The robots were aligned so that they covered almost the whole environment. Typical




Figure 2: Typical data sets obtained with three robots tracking a person in a home environment.



Figure 1: Pioneer I robot used to record the data (left) and Person moving in the environment (right).

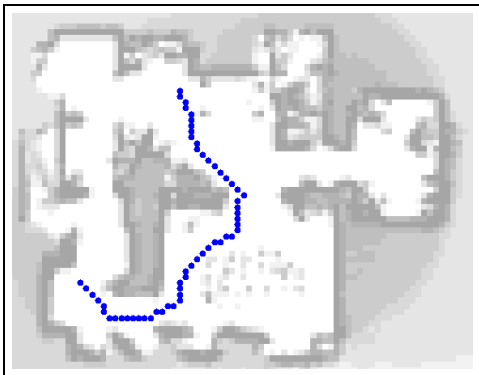


Figure 3: Trajectory of a person extracted from the laser data.

range data obtained during the data acquisition phase are depicted in Figure 2.

To determine the trajectories that are the input to our algorithm we first extract the position of the person in the range scans. We locate changes in consecutive laser-range scans and use local minima in the distance histograms of the range scans. In a second step we identify resting places and perform a segmentation of the data into different slices in which the person moves. Furthermore, we smooth the data to filter out measurement noise. Finally, we compute the trajectories, i.e. the sequence of cells covered by the person during that motion. A typical result of this process is shown in Figure 3.

4 Experimental Results

To evaluate the capabilities of our approach, we performed extensive experiments in a domestic residence as well as in an office environment. Maps of these environments are depicted in Figures 3 and 7. The first set of experiments described here demonstrates the ability of our approach to learn different motion patterns from a set of trajectories. The goal of the second set of experiments is to analyze the classification performance of learned models.

4.1 Application of EM

In the first experiment, we applied our approach to learn a motion model for 42 trajectories recorded in a home environment. Figure 4 shows for one experiment the expectations that were computed in each round of the EM given 14 possible motion behaviors. In this particular experiment, we have exactly three trajectories for each motion pattern. Each column in the pictures contains the expectations $E[c_{i1} | \theta, d], \dots, E[c_{iM} | \theta, d]$ for every trajectory d_i . To enhance the readability, we grouped the examples belonging to the same motion pattern so that they appear as blocks of three trajectories in Figure 4.

Since a uniform distribution of $E[c_{im} | \theta, d]$ represents a local maximum in the log-likelihood space, the EM-algorithm immediately gets stuck if we start with a uniform distribution. We therefore initialize the expectation with a unimodal distribution for each trajectory, i.e., for each d_i the expectations $E[c_{i1} | \theta, d], \dots, E[c_{iM} | \theta, d]$ form a distribution with a unique peak. The location of the mode, however, is chosen randomly.

The topmost image of Figure 4 depicts the initial expectation generated according to the scheme described above. In step 3 the EM has converged to a local maximum in the log-likelihood space. As can be seen from the figure, three trajectories are assigned to two different model components with the same likelihood. Moreover, there are two categories of trajectories that are assigned to the same motion pattern. In step 4 our algorithm therefore removes one of the duplicate models and introduces a new one to which it assigns the trajectory 12 which has the lowest likelihood given the current model θ . After the next iteration, the system has converged to a state

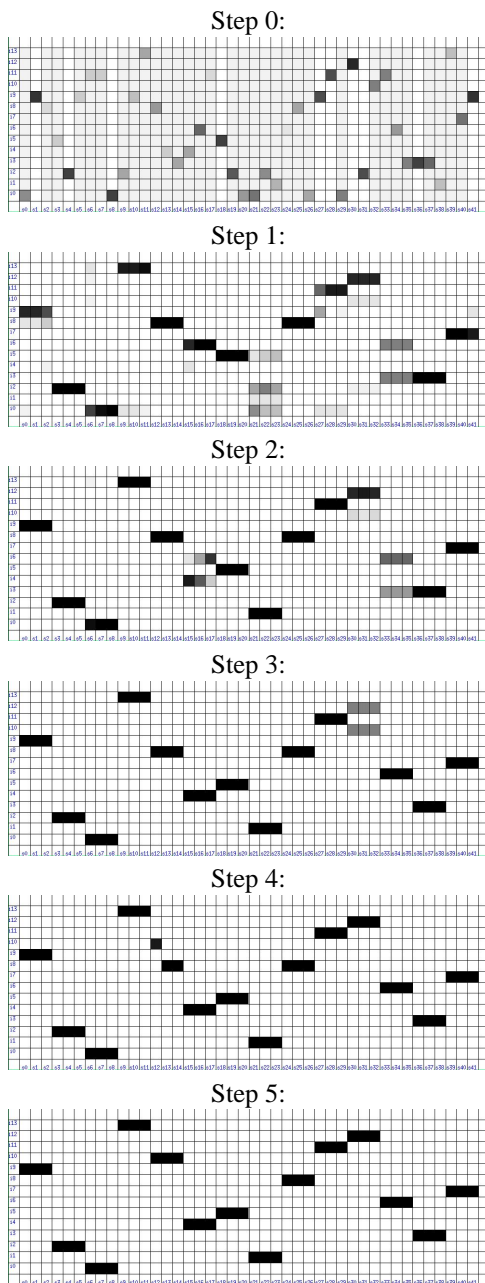


Figure 4: Expectations $E[c_{mn}]$ computed in the different iterations of the EM-algorithm.

in which all trajectories are correctly assigned to the different motion patterns.

To illustrate that our algorithm has correctly clustered the trajectories Figure 5 shows trajectories of two different classes of motion behaviors after the convergence of the EM.

4.2 Predicting Trajectories

To evaluate the capability of our learned models to predict human motions we performed a series of experiments. In each experiment we randomly chose starting fractions of test trajectories and counted the cases in which our model correctly pre-

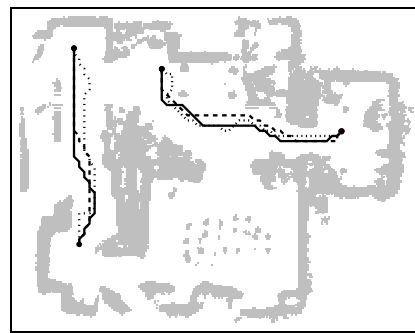


Figure 5: Trajectories of two different classes of motion behaviors.

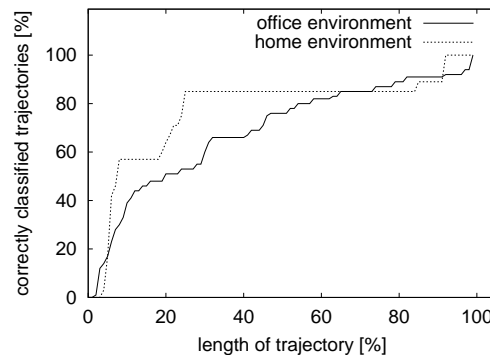


Figure 6: Likelihood of the correct motion behavior after observing fractions of trajectories.

dicts the correct motion behavior. Figure 6 shows in percent the number of correctly predicted motion behaviors depending on the length of the observed trajectory. As can be seen from the figure, the classification results are quite good and our approach yields models allowing a mobile robot to reliably identify the correct motion pattern. If the robot observes 30% of a trajectory in the home environment, then the motion behavior with the highest probability corresponds to the correct motion behavior in over 80% of all cases. The performance in the office environment is not as good as in the home environment, which is because many behaviors have larger parts in common in this environment.

Figure 6 illustrates for one trajectory of the person in the office environment the evolution of the set of possible motion behaviors. Shown in grey are the means of four different motion patterns. The black line corresponds to the trajectory of the person which was observed for the first time at the position labeled S. In the beginning there are four possible motion behaviors (W, B, D, M) to which the trajectory might belong. When location 1 is reached the motion behavior W can be eliminated from the set of hypotheses because the corresponding likelihood gets too low. Thus, even if the system is not able to uniquely determine the intended goal location, it can already predict that the person will follow the corridor during the next steps. When the person reaches location 2 the system can also exclude the motion behavior B. Finally, when the

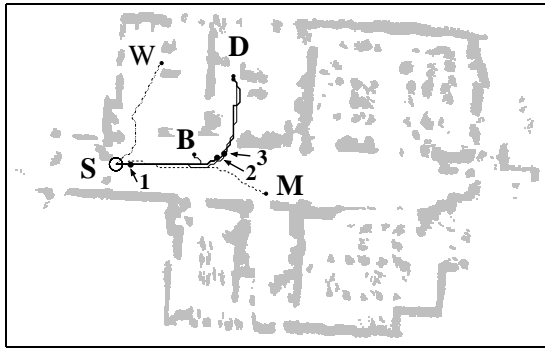


Figure 7: Motion patterns and trajectory of a person.

person reaches position 3, C becomes unlikely and D becomes the most probable motion behavior. This illustrates, that the results of the prediction are useful even in situations in which there are ambiguities about the actual intention of the person.

5 Conclusions

In this paper we presented a method for learning motion behaviors of persons in indoor environments. To cluster similar behaviors into single motion patterns, we apply the popular expectation maximization algorithm. The output of our algorithm is a collection of motion patterns. We furthermore described how to use the resulting models to predict the motions of persons in the vicinity of the robot.

Our approach has been implemented and applied to range data recorded with mobile robots equipped with laser sensors. Special techniques allow the EM-algorithm to overcome local maxima in the likelihood-space which frequently occur in this application. In practical experiments we demonstrated that our method is able to learn typical motion behaviors of a person in a domestic residence as well as in an office building. Based on the resulting motion patterns our system can reliably predict the motions of persons based on observations made by the robot.

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] H. Bui, S. Venkatesh, and G. West. Tracking and surveillance in wide-area spatial environments using the Abstract Hidden Markov Model. *Intl. J. of Pattern Rec. and AI*, 2001.

[3] W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lake-meyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1-2), 1999.

[4] 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.

[5] Z. Ghahramani and S. Roweis. Learning nonlinear stochastic dynamics using the generalized em algorithm. In *Advances in Neural Information Processing Systems (NIPS)*. MIT Press, 1998.

[6] M. Kasper, G. Fricke, K. Stuernagel, and E. von Puttkamer. A behavior-based mobile robot architecture for learning from demonstration. *Journal of Robotics and Automation*, 34(2-3):153–164, 2001.

[7] S. King and C. Weiman. Helpmate autonomous mobile robot navigation system. In *Proceedings of the SPIE Conference on Mobile Robots*, pages 190–198, Boston, MA, November 1990. Volume 2352.

[8] F. Kruse, E. und Wahl. Camera-based monitoring system for mobile robot guidance. In *Proc. of the International Conference on Intelligent Robots and Systems (IROS)*, 1998.

[9] G. Lacey and K Dawson-Howe. The application of robotics to a mobility aid for the elderly blind. *Journal of Robotics and Autonomous Systems (RAS)*, 23:245–252, 1998.

[10] S. M. Lavalle, H. H. Gonzalez-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.

[11] G.J. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. Wiley Series in Probability and Statistics, 1997.

[12] T. M. Mitchell. *Machine Learning*. McGraw-Hill, 1997.

[13] N. Roy, G. Baltus, D. Fox, F. Gemperle, J. Goetz, T. Hirsch, D. Magaritis, M. Montemerlo, J. Pineau, Schulte J., and S. Thrun. Towards personal service robots for the elderly. In *Proc. of the Workshop on Interactive Robotics and Entertainment*, 2000.

[14] C. Schaeffer and T. May. Care-o-bot - a system for assisting elderly or disabled persons in home environments. In *Assistive technology on the threshold of the new millenium*. IOS Press, Amsterdam, 1999.

[15] R.D. Schraft and G. Schmierer. *Serviceroboter*. Springer Verlag, 1998. In German.

[16] D. Schulz, W. Burgard, D. Fox, and A.B. Cremers. Tracking multiple moving targets with a mobile robot using particle filters and statistical data association. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2001.

[17] S. Tadokoro, M. Hayashi, Y. Manabe, Y. Nakami, and T. Takamori. On motion planning of mobile robots which co-exist and cooperate with human. In *Proc. of the International Conference on Intelligent Robots and Systems (IROS)*, 1995.

[18] Q. Zhu. Hidden Markov model for dynamic obstacle avoidance of mobile robot navigation. *IEEE Transactions on Robotics and Automation*, 7(3), 1991.