

Adapting Navigation Strategies Using Motions Patterns of People

Maren Bennewitz and Wolfram Burgard
Department of Computer Science
University of Freiburg, Germany

Sebastian Thrun
School of Computer Science
Carnegie Mellon University, PA, USA

Abstract—As people move through their environments, they do not move randomly. Instead, they are often engaged in typical motion patterns, related to specific locations they might be interested in approaching. In this paper we propose a method for adapting the behavior of a mobile robot according to the activities of the people in its surrounding. Our approach uses learned models of people’s motion behaviors. Whenever the robot detects a person it computes a probabilistic estimate about which motion pattern the person might be engaged in. During path planning it then uses this belief to improve its navigation behavior. In different practical experiments carried out on a real robot we demonstrate that our approach allows a robot to quickly adapt its navigation plans according to the activities of the persons in its surrounding. We also present experiments illustrating that our approach provides a better behavior than a standard reactive collision avoidance system.

I. 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. Within this paper we consider the problem of how knowledge about typical motion patterns of people can be utilized to improve the navigation behavior of the robot. In particular, we are interested in predicting the motions of persons and instructing the robot to choose appropriate detours so that the risk of interferences with persons is minimized.

As an example, consider the situations illustrated in Figure 1. In the left image a robot is entering a corridor and moving to the left approaching its designated target location. At the same time, a person is walking from left to right in the same corridor. In this particular situation, the robot needs to be able to detect the person and to predict its future actions in order to prevent interfering with it. A similar situation is depicted in the right image. Here, the robot is standing in a doorway and a person, that




Fig. 1. Situations in which a robot interferes with a person. In both cases the knowledge that the person’s path will lead through the position of the robot would help the robot to avoid this conflict.

wants to enter the room, is approaching the robot. Again, if the robot fails to detect the intention of the person and to react appropriately, the person cannot enter the room immediately.

The goal of this paper is to introduce a technique that allows a mobile robot to predict future motions of persons and to incorporate this knowledge into its navigation plans. In particular, we describe a probabilistic technique to determine potential motion patterns of persons detected by a robot. The knowledge about potential intentions is then used to plan the actions of the robot in its configuration time-space. This way, the robot is able to avoid interfering with persons by choosing trajectories that stay away from the predicted paths of the persons in its surrounding.

This paper is organized as follows. After discussing related work in the following section, we briefly describe the representation of the motion patterns. In Section IV we explain how the learned motion models can be integrated into a path planning system. After briefly describing our laser-based implementation, we present several experiments in Section VI. There, we focus on the ability of our approach to improve the behavior of mobile robots by adapting their navigation plans when observing persons.

II. RELATED WORK

Recently, a variety of service robots have been developed that have been designed to operate in populated environments. These robots, for example, have been deployed in hospitals [8], museums [5], office buildings [1], and department stores [6], where they perform various services, e.g., deliver, educate, entertain [17] or assist people [16], [11], [15]. 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 example, the techniques presented by Schulz et al. and Prassler et al. [18], [9] are designed to track multiple persons in the vicinity of a robot. The technique described by Bui et al. [4] uses an Abstract Hidden-Markov-Model to learn and predict motions of a person. The authors, however, do not explain how to cluster different motions into motion patterns and how to exploit the learning results to control a mobile robot. Recently, we have developed a technique to learn motion patterns of persons [3] based on typical trajectories of the persons in their environment. In this paper, we describe how these learned motion patterns can be used to improve the navigation behavior of a robot.

In the past, different techniques were introduced to adapt the behavior of a robot according to the knowledge about the actions of people in its surrounding. For example, the approaches presented by Tadokoro et al. [20] and Zhu [21] apply Hidden-Markov-Models to predict the motions of moving obstacles in the environment of a robot. Since these approaches do not learn typical motion patterns, they can only predict short-term motions and not complete trajectories. The system described by Kruse and Wahl [10] uses cameras to track people and to learn where persons typically walk. In contrast to our system, Kruse and Wahl do not predict the motion of an observed person. When the path of the robot is blocked the robot stops and replanning is invoked, treating the unforeseen obstacle as a static obstacle. Kasper et al. [7] presented an approach to improve the behavior of a robot by following the activities of a teacher. Latombe et al. [12] developed 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. The major difference between the latter two approaches and our technique lies in the different evaluation functions. Whereas Kasper et al. seek to optimize the navigation skills, Latombe et al. generate actions to maximize the visibility of a moving object. The approach presented in this paper, in contrast, has the goal to minimize the risk of interfering with persons given the knowledge about typical motion patterns.

III. MOTION PATTERNS

Our approach uses motion patterns of people that are learned using the technique described in [3]. The key idea of this method is to cluster trajectories of persons into motion pattern using the EM-algorithm. The output is a number M of different types of motion patterns $\theta = \{\theta_1, \dots, \theta_M\}$ a person might exhibit in its natural environment. A motion pattern, denoted θ_m with $1 \leq m \leq M$, is represented by K Gaussian distributions $\theta_m^{[k]} = N(\mu_m^{[k]}, \sigma)$ with mean $\mu_m^{[k]}$ and fixed standard deviation σ for all m and k . Each such Gaussian specifies for each data point x^i and each $\theta_m^{[k]}$ the likelihood $p(x^i | \theta_m^{[k]})$ that the person is at location x^i given that step i of the trajectory x corresponds to step k of motion pattern m :

$$p(x^i | \theta_m^{[k]}) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2\sigma^2} \|x^i - \mu_m^{[k]}\|^2}. \quad (1)$$

As in [3], we assume that all motion patterns are of the same length K , which can be achieved in a straightforward way by linear interpolation.

In the remainder of this section we describe how the learned motion patterns can be used to predict motions of surrounding people. Suppose the robot observes a sequence $z = \{z^1, z^2, \dots, z^R\}$ of positions of a person. What we are interested in is a distribution which gives us for each motion pattern θ_m the probability $p(\theta_m | z)$ that

the person is engaged in θ_m given z . According to Bayes' rule, this corresponds to

$$p(\theta_m | z) = \alpha p(z | \theta_m) p(\theta_m). \quad (2)$$

Here, $p(z | \theta_m)$ is the likelihood of the data given θ_m , $p(\theta_m)$ is the prior for θ_m , and α is a normalizer ensuring that the left-hand side sums up to one over all θ_m .

It remains to describe how $p(z | \theta_m)$ is computed. Unfortunately, z does not necessarily start at the initial position of the corresponding motion pattern. Suppose $\theta_m^{[k]}$ with $1 \leq k \leq K$ is the position in θ_m the first observed position z^1 corresponds to. Furthermore, suppose $\theta_m^{[k']}$ with $k \leq k' \leq K$ is the position of θ_m the final observation z^R of z corresponds to. Since both k and k' are unknown, we apply the law of total probability and compute $p(z | \theta_m)$ by summing over all possible combinations of k and k' :

$$p(z | \theta_m) = \sum_{k=1}^K \sum_{k'=k}^K p(z | \theta_m, k, k') p(k, k'). \quad (3)$$

The prior probability $p(k, k')$ depends on the velocity and on the length of the given segment on the motion pattern. To get $p(z | \theta_m, k, k')$ we compute the product of the likelihoods of each observation z^r in z given that it starts at k and ends at k' :

$$p(z | \theta_m, k, k') = \prod_{r=1}^R p(z^r | \theta_m, k, k') \quad (4)$$

$$= \prod_{r=1}^R p(z^r | \theta_m^{[f(r,k,k')]}]) \quad (5)$$

where

$$f(r, k, k') = \frac{k' - k}{R - 1} r + \frac{kR - k'}{R - 1} \quad (6)$$

realizes a linear mapping of the individual observations z^1, \dots, z^R to the components $\theta_m^{[k]}, \dots, \theta_m^{[k']}$ of θ_m .

IV. INTEGRATING PREDICTED MOTIONS OF PERSONS INTO PATH PLANNING

In the previous section we described how to estimate the probability that a person is engaged in a motion pattern θ_m given an observation sequence z . Within this section we now focus on the question of how the robot can exploit this information to improve its navigation behavior. In particular, we want to focus on the question how the belief of the robot about the intentions of surrounding persons can be considered during the path planning process.

Our robot applies the A^* procedure for path planning and searches for the minimum-cost path in its three dimensional configuration time-space. The environment is represented as a static occupancy grid map [14]. Each cell $\langle x, y \rangle$ of this grid stores the probability $P(\text{occ}_{x,y})$ that the corresponding area in the environment is occupied. As

in [2] the cost for traversing a cell $\langle x, y \rangle$ is proportional to its occupancy probability $P(occ_{x,y})$. To avoid that paths lead through walls etc. we apply a threshold function $\gamma(x)$ which is infinite if x exceeds 0.8 and x elsewhere. As a heuristics we use the value function obtained by a deterministic value iteration in the static 2D map. This allows the robot to quickly re-plan its path whenever new measurements have arrived and the belief of the intended trajectories of the persons has changed.

To incorporate the robot's belief about future trajectories of the persons, we additionally discount a cell $\langle x, y \rangle$ according to the probability, that one of the persons covers $\langle x, y \rangle$ at a given time t . Suppose our robot has observed L persons and suppose $p(cov_{x,y,t} | z_l)$ is the probability that a person l covers $\langle x, y \rangle$ at time t given the observations z_l corresponding to this person. If we consider the individual persons independently, we can then compute the costs $C_{cov}(x, y, t)$ introduced by the fact that $\langle x, y \rangle$ might be covered at time t as follows:

$$C_{cov}(x, y, t) = \sum_{l=1}^L p(cov_{x,y,t} | z_l) \quad (7)$$

$$= \sum_{l=1}^L \sum_{m=1}^M \sum_{k=1}^K \sum_{k'=1}^K \left(p(cov_{x,y,t} | \theta_m, k, k', z_l) \cdot p(\theta_m, k, k' | z_l) \right) \quad (8)$$

It remains to describe how the individual terms $p(cov_{x,y,t} | \theta_m, k, k', z_l)$ and $p(\theta_m, k, k' | z_l)$ are computed. For the latter term we have

$$p(\theta_m, k, k' | z_l) = \alpha p(z_l | \theta_m, k, k') p(\theta_m, k, k') \quad (9)$$

$$= \eta p(z_l | \theta_m, k, k') p(\theta_m) p(k, k') \quad (10)$$

where η is another normalizer. Finally, we need to describe how we actually compute the probability $p(cov_{x,y,t} | \theta_m, k, k', z_l)$ that a person engaged in motion pattern θ_m will cover $\langle x, y \rangle$ at time t given the observations z_l and given that z_l starts at $\theta_m^{[k]}$ and ends at $\theta_m^{[k']}$. In our current system, we are using a Gaussian distribution to represent the uncertainty about the position of the person at time step t . The mean of this Gaussian is computed as that point on θ_m which has the distance vt from the latest observed position t' , where v is the velocity of the person in the observations z_l . Thus, we predict the motion of the person starting from location t' according to the average velocity v and the trajectory given by θ_m .

The overall cost for the robot to traverse a cell $\langle x, y \rangle$ at time t is then computed as:

$$C(x, y, t) = \gamma(p(occ_{x,y})) + C_{cov}(x, y, t) \quad (11)$$

At this point it is worth noting that our approach can be used to predict the intentions of multiple persons and that it can even deal with persons not engaged in any of the learned motion patterns. If a person's behavior cannot be associated well to any of the known models, the belief $p(\theta_m | z)$ will be uniformly distributed which introduces higher costs for fields on learned motion patterns.

V. DETECTING AND TRACKING PEOPLE

To apply the technique described above, a robot must be able to detect persons in its vicinity and to keep track of them. Similar to [18], [13], our current system extracts features out of range scans and considers changes in consecutive scans to identify moving people. To keep track of a person we use a Kalman filter. The state x of a person is represented by a vector $[x, y, \delta x, \delta y]'$. Whereas x and y represent the position of the person, the terms δx and δy represent the velocity of the person in x - and y -direction. Accordingly, the prediction is carried out by the equation:

$$x_{r+1}^- = \begin{bmatrix} 1 & 0 & t_r & 0 \\ 0 & 1 & 0 & t_r \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} x_r \quad (12)$$

where t_r is the time elapsed between the measurement y_{r+1} and y_r . Usually, sensors only give the position of an object. Since the sensors we are using generally do not provide the velocities δx and δy , which are also part of our state space, the measurement matrix projects onto the first two components of the state space. Accordingly, the predicted measurement at step $r + 1$ is:

$$y_{r+1}^- = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} x_{r+1}^- \quad (13)$$

The positions z^i of a person that are input to our path planning routine, are computed out of each x_i by using the first two components of the state space. To deal with multiple persons we use independent Kalman filters to keep track of them. To solve the data association problem, we apply a nearest neighbor approach, i.e. we update a filter using that observation y_{r+1} that is most closely to y_{r+1}^- . New filters are introduced for observations from which all predicted observations are too far away. Furthermore, filters are removed if no updates can be made to a filter for several update steps.

VI. EXPERIMENTAL RESULTS

To evaluate the capabilities of our approach, we performed extensive experiments. The experiments demonstrate that our robot is able to use the learned models to classify observed trajectories, predict motions of persons and adapt its own behavior accordingly. All experiments were carried out using our B21r platform Albert in the




Fig. 2. Albert moves into a doorway to let the person pass by.




Fig. 3. Albert moves forward and waits until the likelihood of interfering with the person is low enough.

corridor environment of the department of Computer Science at the University of Freiburg. Albert is equipped with a laser-range finder which is used in the experiments reported here for people tracking and obstacle avoidance.

A. Planning Detours

The first experiment is designed to demonstrate that our approach allows a mobile robot to reliably predict the possible trajectories of persons and to appropriately incorporate this information into its motion plans. The task of the robot was to travel along the corridor of our building. At the same time, a person walking in the opposite direction was approaching the robot. Figure 2 shows the initial position of the robot and the position of the person when it was discovered by the robot for the first time. Given the first two estimates about the person's position and the existing motion patterns, possible trajectories lead to the locations D, M, E and W which are also shown in Figure 2. Obviously, all the corresponding motion patterns lead through the corridor so that the robot was likely to interfere with the person. Accordingly, the cost-optimal action for the robot was to drive into the doorway to the right in front of it and to wait there until the person eventually had passed by.

Figure 2 also shows the whole trajectory of the person (dashed line) as well as the trajectory of the robot (solid line). As can be seen from the figure, the person went to the location E and the robot continued to move towards its designated goal point after the person had walked by.

Figure 3 shows a similar situation. Again, the robot had been driving in the corridor when it realized a person approaching it. According to the observations, the most likely motion patterns of the learned model lead to the locations D, M, L, and S. However, since the prior probabilities of the motion patterns with target locations L and S were very low, the resulting costs introduced to the configuration time-space did not prevent the robot from driving further. Since the other possible intentions leading the person to D and M had a sufficiently high prior, the system stopped in the middle of the corridor and waited until the person entered room M in this case. Figure 4

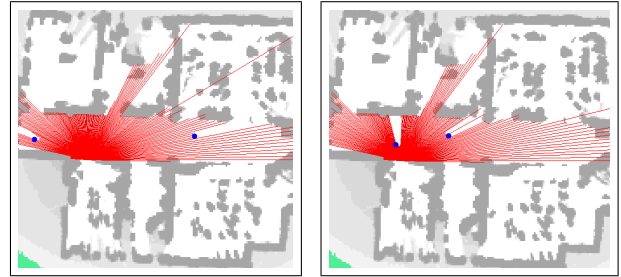


Fig. 5. Tracking two persons using Kalman filters.

shows images of the robot and the person taken during this experiment. As can be seen, the robot successfully avoided that the person had to take a detour in order to arrive in the target location.

B. Multiple Persons

Figure 6 shows a situation in which Albert was about to leave a room while two persons were walking along the corridor and approaching the robot from either direction. Figure 5 shows laser-range scans and the estimated positions of the persons at different time steps. The trajectories of the persons are depicted as solid lines in Figure 6. Since Albert was not able to leave the room before person P_1 had walked by it stayed in the doorway. After P_1 had passed the doorway the behavior of the robot was mainly influenced by the intention of P_2 . Since P_2 continued walking along the corridor Albert waited until also P_2 had passed the doorway. If, in contrast, P_2 had walked into one of the offices before passing the robot or P_2 had moved at a lower speed, Albert would have started moving immediately after P_1 had unblocked the robot's path. The possible intentions are also depicted in Figure 6.

C. Comparison to a Reactive Collision Avoidance System

Additionally, we compared our approach to a reactive collision avoidance [19]. Figure 7 shows the behavior of the robot in the same situation as depicted in Figure 2. As can be seen from the figure, the robot immediately started to move towards its target location. As soon as the




Fig. 4. Albert moves forward and waits until the likelihood of colliding with the person is low enough.




Fig. 6. Albert observes two persons and waits in the doorway until they have passed it.




Fig. 8. Albert moves aside in order to let a person pass by.




Fig. 7. If the robot does not use our predictive planning system and solely relies on a reactive collision avoidance system in the same situation as depicted in Figure 2, it will get pretty close to a person and will force it to slow down.




Fig. 9. Albert moves away from a doorway in order to let a person enter the corresponding room.

person was close by, the robot tried to move around it. Please note that in this case the person had to reduce its speed and let the robot move aside before it could go on. This illustrates, that our system provides a much better behavior in dynamic environments than purely reactive collision avoidance systems.

D. Giving Space to Persons

The final experiment described here is designed to illustrate that our technique can also be used to improve the behavior of the robot even in situations in which the robot is not performing a navigation task. In this particular situation (see Figure 8) the robot had no goal point and rested in a doorway waiting for instructions. Then it realized that a person was approaching from the left. According to the learned motion patterns, Albert

inferred a high chance that the person would enter the room through the doorway that was blocked by it. The cost-optimal action according to the path planner was to drive to the middle of the corridor in order to give space. Figure 9 shows two images taken during this experiment.

VII. CONCLUSIONS

In this paper we presented a method for adapting motion strategies of a mobile robot according to the activities of surrounding people. Our approach uses motion patterns learned out of real data and exploits these patterns to predict the motions of persons sensed by the robot. To compute cost-optimal paths that minimize the risk of interfering with a person we consider the configuration time-space of the robot. The integration of the estimated motion patterns of the persons into the path planning allows the robot to adapt its behavior more appropriately and at an earlier stage than purely reactive approaches.

Our technique has been implemented and applied to data recorded with mobile robots equipped with laser-

range sensors. The current implementation is highly efficient and allows to quickly react to its sensory input. In different experiments we demonstrated that the behavior of a mobile robot can be improved by predicting the motions of surrounding people. They furthermore illustrated advantages over standard reactive systems.

ACKNOWLEDGMENTS

We would like to thank Cyrill Stachniss for his efforts in carrying out the experiments with our mobile robot.

This work reported here has partly been supported by the IST Programme of Commission of the European Communities under contract number IST-2000-29456.

VIII. 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 Int. Joint Conference on Artificial Intelligence (IJCAI)*, 1997.
- [2] M. Bennewitz, W. Burgard, and S. Thrun. Exploiting constraints during prioritized path planning for teams of mobile robots. In *Proc. of the Int. Conference on Intelligent Robots and Systems (IROS)*, 2001.
- [3] M. Bennewitz, W. Burgard, and S. Thrun. Using EM to learn motion behaviors of persons with mobile robots. In *Proc. of the Int. Conference on Intelligent Robots and Systems (IROS)*, 2002.
- [4] 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.
- [5] W. Burgard, A.B. 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.
- [6] H. Endres, W. Feiten, and G. Lawitzky. Field test of a navigation system: Autonomous cleaning in supermarkets. In *Proc. of the Int. Conference on Robotics & Automation (ICRA)*, 1998.
- [7] 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.
- [8] 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, November 1990. Volume 2352.
- [9] 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 Int. Conference on Robotics & Automation (ICRA)*, 2001.
- [10] E. Kruse and F. Wahl. Camera-based monitoring system for mobile robot guidance. In *Proc. of the Int. Conference on Intelligent Robots and Systems (IROS)*, 1998.
- [11] 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.
- [12] 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 Int. Conference on Robotics & Automation (ICRA)*, 1997.
- [13] M. Montemerlo, S. Thrun, and W. Whittaker. Conditional particle filters for simultaneous mobile robot localization and people-tracking. In *Proc. of the IEEE Int. Conference on Robotics & Automation (ICRA)*, 2002.
- [14] H.P. Moravec and A.E. Elfes. High resolution maps from wide angle sonar. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 116–121, 1985.
- [15] 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.
- [16] 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.
- [17] R.D. Schraft and G. Schmierer. *Service Robots*. Springer Verlag, 1998.
- [18] 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 Int. Conference on Robotics & Automation (ICRA)*, 2001.
- [19] C. Stachniss and W. Burgard. An integrated approach to goal-directed obstacle avoidance under dynamic constraints for dynamic environments. In *Proc. of the Int. Conference on Intelligent Robots and Systems (IROS)*, 2002.
- [20] S. Tadokoro, M. Hayashi, Y. Manabe, Y. Nakami, and T. Takamori. On motion planning of mobile robots which coexist and cooperate with human. In *Proc. of the Int. Conference on Intelligent Robots and Systems (IROS)*, 1995.
- [21] Q. Zhu. Hidden Markov model for dynamic obstacle avoidance of mobile robot navigation. *IEEE Transactions on Robotics and Automation*, 7(3), 1991.