Socially Inspired Motion Planning for Mobile Robots in Populated Environments

Jörg Müller Cyrill Stachniss Kai O. Arras Wolfram Burgard

***Abstract*— Nowadays, robots often operate in environments that they share with humans. The ability to act similar to humans is an important prerequisite for the social acceptance of robots. In this paper, we consider the problem of navigation in populated environments. We present a path planning algorithm that enables robots to move efficiently and smoothly with groups of people by selecting those individuals that move towards the robot's desired goal. Our technique is based on a people tracking system in combination with an iterative A* planner. The approach iteratively finds both, a path and a partition of the set of surrounding people into obstacles and subjects to follow. In the absence of people, the optimal solution is still found by the A* planner. The approach has been implemented and tested on a real mobile robot in populated environments. Experiments illustrate that the robot is able to move with groups of people resulting in a more human-like way of navigation among people.**

I. Introduction

Whenever humans navigate in populated environments, they often move in groups. This walking strategy leads to fewer detours and avoids, for example, slightly slower walking people. It can be expected that the ability to adapt to moving people improves the navigation capabilities of a robot and leads to a more human-like behavior. This is opposed to most approaches to autonomous robot navigation that seek to find the shortest trajectory to the desired goal or minimizes the risk of collisions with static or dynamic obstacles. None of these approaches (see next section) aims to guide a robot through a crowded environments similar to the way humans do.

This paper presents a novel approach to autonomous navigation that allows a robot to move similar to humans. In populated environments, our approach seeks to move *with* groups of people or individuals as long as they move towards the robot's goal. Our technique is based on a people tracking system in combination with an iterative A* planner. In case the robot encounters an obstacle-free trajectory, the methods guides the robot on the shortest possible trajectory. Our method has been implemented and successfully tested on a real robot. Applications include robots that work in museums, hospitals, public spaces, trade fairs as well as autonomous wheelchairs.

II. Related Work

Whenever robots are deployed in populated environments, the ability to perceive people and react according to their actions improves the service and robustness of a mobile robot. In the past, several researcher addressed the problem of detecting and tracking humans over time. Montemerlo and Thrun [12] proposed the SLAP framework that simultaneously tracks the poses of people and localizes the robot in the environment. Schulz *et al.* [16] proposed an approach that is able to robustly track multiple people. To address the data association problem, they applied joint probabilistic data association filters in combination with a particle filter.

Given the knowledge about locations of people, the robot can adapt its behavior and interact with the individuals. For example, the museum tour guide RHINO [5] gave successful tours in different museums. In case of crowded spaces it needed to traverse, the authors used speech synthesis to request people to let the robot path through. The tour guide robot Minerva [20] used a similar system while operating in the Smithsonian museum of American history. During Expo.02, the robots deployed by Siegwart *et al.* [17] used a combination of global planning and local collision avoidance for navigation. For the same exposition project, Jensen [9] distinguishes between static, semi-static, and dynamic objects. This allows to more robustly track the individual objects in the surroundings of the robot. Even so these three robots operated in highly populated environments, their path planning approaches did not adapt to the motion and estimated movements of observed people.

Bennewitz [3] proposed an approach to learn typical motion patterns of people and to adapt the motion planning according to the predicted motions. This allows for more accurate motion prediction but requires to learn models of the individual people. The learning of such patterns is performed by an EM style algorithm [4]. In contrast to this, our approach operates without explicitely learned models. To integrate knowledge about people in the path planning process, Foka and Trahanias [7] presented an approach that uses a partially observable Markov decision processes (POMDPs) for predicting the motion of people. In this way, their approach allows a mobile robot to plan trajectories that are less likely to interfere with people.

The approach of Tadokoro *et al.* [19] also applies probabilistic techniques to predict the motion of people based on external sensors. They apply a genetic algorithm for path planning that seeks to minimize the length of the trajectory and the risk of collisions. Philippsen *et al.* [13] uses an extension of ICP that combines motion detection from a mobile platform with position estimation. This information is used to estimate a traversability risk function that unifies dynamic and static obstacles. The planning system uses this risk function to generate smooth paths that trade off collision risk versus detours.

Dealing with dynamic environments is also an important topic for autonomous or semi-autonomous wheelchairs. The system of Prassler *et al.* [14] perform short term motion prediction of dynamic obstacles in order to ensure safe navigation without collisions using some heuristics.

All authors are members of the University of Freiburg, Department of Computer Science, D-79110 Freiburg, Germany

All techniques presented so far estimated the location and motion of people in the environment. Based on this knowledge, these approaches aim to plan short trajectories that minimize the risk of collision. This results in avoiding regions close to people and prevents the robot from adapting its motion to the dynamics or the humans. In contrast to this, our approach presented in this paper seeks to adapt its motion to the way humans move. Instead of driving as fast as possible on the preferably shortest trajectory that is collision free, our robot is able to move with groups of people and still follows a person even if it would be possible to overtake. We achieve this by an iterative planning approach that seeks for people that appear to move to the same target location than the robot.

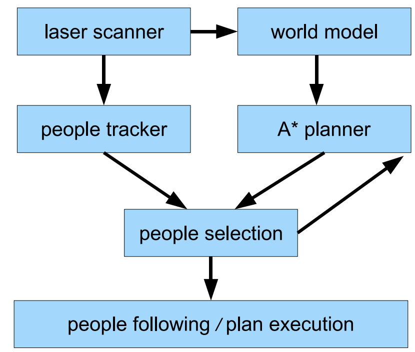

Fig. 1. The information flow of the individual components.

III. Laser-based People Tracking

A prerequisite for a socially inspired approach to navigation is the ability of the robot to reliably detect and track people in its surroundings. This section describes the Kalman filter-based multi-target tracker that is utilized to detect and track people. We briefly go through the tracking cycle. For the details of Kalman filtering and target tracking the reader is referred to [2].

State prediction. A people track is represented as $\mathbf{x} = (x, y, v_x, v_y)$ where x and y are the position and v_x and v_y the x and y velocity components of the person. With this state representation, new tracks can be properly initialized with $v_x = v_y = 0$. For motion prediction, a constant velocity model is employed.

Measurement prediction. As the x- and y-coordinates of a track are directly observable, the 2×4 measurement matrix H is formed by the 2×2 identity matrix in x and y and the 2×2 zero matrix in v_x and v_y.

Observation. In the observation step, people are detected in the laser range data. The problem can be seen as a classification problem that consists in finding those laser beams that correspond to people and to discard other beams. Typically, hand designed classifiers have been employed for this task. A popular approach is to extract people by detecting moving blobs that appear as local minima in the laser range scans. Neither the selection of features nor their thresholds are learned or determined other than by manual design and hand-tuning [10, 6, 15, 16, 21]. This motivates the application of a supervised learning technique. In particular, our approach applies AdaBoost to train a strong classifier from simple features of groups of neighboring beams corresponding to legs in range data [1]. AdaBoost takes a labeled training set and a possibly large vocabulary of features that may or may not be appropriate for the given classification task. The method then creates a classifier by selecting the most informative features and finding the optimal thresholds given the training data. Experiments show that for people detection, the AdaBoost classifier is superior to a manually designed classifier and is therefore also used in this work.

Data association. For data association, we employ a nearest neighbor standard filter that considers tracks separately. Although this is a very simplistic approach to data association, the performance was sufficient for our purposes. Our current extension of this tracker integrates data association decisions over time using multiple hypotheses.

Estimation. Given that both, the state and measurement prediction models are linear, a (non-extended) Kalman filter as the optimal estimator under the Gaussian assumption can be employed.

The tracker operates in the world reference frame, i.e. the track position and velocities are represented in world coordinates. Observations acquired in the sensor frame are transformed into the world frame using the robot's pose known from localization. Uncertainties are transferred from local to global coordinates using first-order error propagation.

In practice, the tracker cycle frequency is much higher than the rate with which new localization updates arrive. We therefore use odometry to estimate the robot's pose as it delivers faster and smoother pose updates. Although this causes the pose and hence the transformed track states to drift, this is only internal to the tracker. Each time when a new localization update is available, the filtered track states are finally transformed into the (drift-free) world reference frame and passed to the planner which also operates in that frame.

IV. Planning and People Selection

This section explain how the detected people are incorporated into our planning approach. Figure 1 shows the information flow between the individual modules. Our approach iterates the A* planner and the people selection technique.

A. Iterative Planning

To decide whether to follow a person and to find the best person to follow, we are looking for people who seem to have similar short-term goals than the robot itself. To determine the people to follow, we iteratively plan a path and select people according to the previously planned path.

For modeling the environment, we use an occupancy grid map which contains only static obstacles and can be obtained with standard mapping techniques [8]. Additionally, we maintain a temporary local map containing the observed (dynamic) obstacles based on the latest n laser observations ($n = 20$). The basic path planning is performed by the A* algorithm on the combination of the global (static) and local occupancy grid map. Thus we plan in the 2-dimensional (x, y)-Space for reason of efficiency. As a fast computable heuristic we use Dijkstra's algorithm on the static map, which is once computed when the goal is set. This leads to fast re-planning capabilities while we guarantee a directed search especially through complicated room- and corridor-combinations [18]. To avoid collisions caused by driving too close to obstacles, we convolve the grid with a Gaussian kernel to give higher costs to cells close to obstacles. This allows the robot to keep a reasonable distance to obstacles but at the same time enables it to navigate through narrow passages.

Since we do not learn prediction of a person's movements like Foka and Trahanias [7] or Bennewitz [3], we predict only the short-term movements of people into our navigation system. Based on the estimated short-term movements, our planner needs to determine if a person is suitable to act as a leader which refers to the fact that the robot follows this person. The robot should select only those people as a potential leader, who move towards its own goal location. To achieve this, we apply an iterative A* planning approach, which is described in the following.

In the initial planning step, we ignore all people and consider only the static obstacles in the environment. Thus, the A* planner reports the best action of the robot given the static objects. The robot then seeks for people that interfere with the planned trajectory. In case there are no people close to the planned path, there is no need for the robot to adapt its trajectory. As a result, in the absence of people, the robot generates the optimal trajectory to the target location.

However, in highly populated environments such as museums or trade fairs people are likely to interfere with the planned path of the robot. In this case, it identifies which person is suitable to become a leader. This is done based on the position and motion direction of the person with respect to the planned trajectory. In its world model, the robot then marks the person as a potential leader or as an obstacle it has to avoid during planning. Then the robot executes the A* algorithm depending on the modified world model, in which the grid cells occupied by people, who are not considered as obstacles, are marked as free. These steps are executed in an iterative fashion until the approach converges to an admissible solution. Here, we call a path admissible if and only if it was planned without considering exactly those people as obstacles, who are suitable to follow based on this path. Obviously, each admissible path is a stable solution of our iterative planning approach. An algorithmic description of our work is depicted in Algorithm 1.

The decision if a person is a potential leader, the robot considers the short-term prediction from the people track in combination with the latest trajectory computed by A*. In detail, we allow a maximum deviation from the plan by 1 m to the right or left hand side as well as a maximum angular deviation of 30 deg.

Since our navigation system frequently re-plans its path to adapt to the changes in the environments, the robot only follows people as long as they move in the same direction than the path guiding the robot to its target location. As as result, the robot moves with people only as long as it is convenient with respect to its own goal.

Figure 2 illustrates an example of our iterative planning in which four people are in the vicinity of the robot and are considered in the planning process. After four iterations, the approach converged to a solution and it is reported as a result of our planning algorithm. After identifying the person or group of people the robot is supposed to follow, one can apply a standard local navigation approach such as potential field, nearness diagram navigation [11], or dynamic window approach [5, 20] that guides the robot to a position which is lying slightly behind the person to follow. We obtained comparably good results using the potential field method for that.

Algorithm 1 Iterative path planning with people

Let $h_i \subseteq P$ the set of people not considered as obstacles in planning step i and $f_i \subseteq P$ the set of people who are suitable to follow based on the path planned in step i.

Input: People $P = \{P_1, \ldots, P_n\}$
Output: Path p and the set of people, who are suitable to follow

```
i = 0
h_0 = P
loop
    Plan path p_i with h_i
    Compute potential leaders f_i based on path p_i
    if f_i = h_i then
        return ⟨p_i, f_i⟩
    end if
    if f_i = h_j for any j < i then
        return "no path found"
    end if
    h_{i+1} = f_i
    i = i + 1
end loop
```

Note that it can happen that our approach does not converge since it generates an intermediate solution that was already generated in a previous iteration smaller $i-1$, where i refers to the current iteration.

In the case in which no admissible path is found, we compute the path which considers all people as obstacles and report this as the final solution. This path is indeed not admissible according to our definition, but it is collision free and thus can be executed by the robot.

B. Blocking

Another feature we added to our navigation approach is the ability to actively approach people who are standing around and thus blocking significant regions in the environment (see Figure 3 for an illustration). Those people often cause the robot to accept detours to reach its goals (if there exists an obstacle free path at all). In practical situations, however, people often move away from narrow passages such as doorways as soon as the robot approaches them. We address this issue by trying to shoo away such people. We compute a path ignoring the person. If this path is significantly shorter than the originally planned path, our system tries to shoo away the person and in this way free the path. To shoo someone away, we approach the person, accelerate shortly in front of them before braking

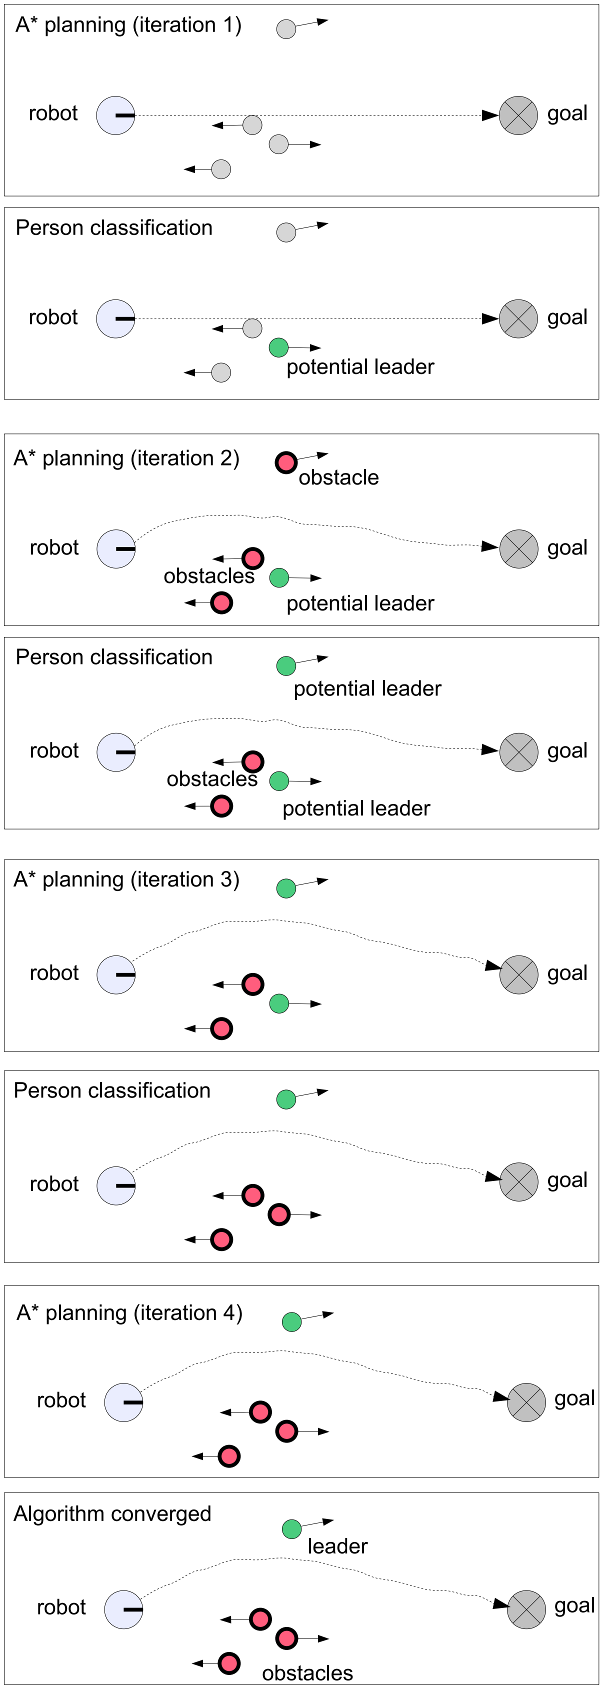

Fig. 2. Example of our iterative algorithm to plan a path and select people to follow. Each odd row depicts the result of the A* planner and each even row the result of the person selection technique. Bold red (dark gray) circles indicate people classified as obstacles while green ones indicate potential leaders.

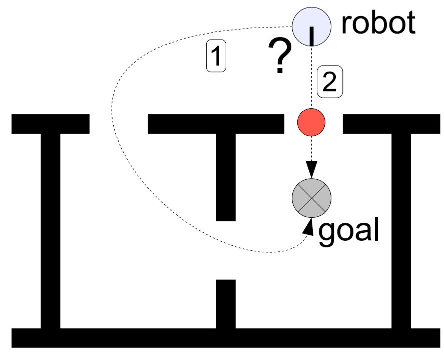

Fig. 3. A person blocks a significantly shorter path causing a detour for the robot. Alternatively, the robot can try to shoo this person away and to move on the shorter path.

again. In most cases, this behavior causes people to intuitively free the path.

V. Experiments

In this section, we present the experiments carried out to evaluate our approach. We used simulated as well as real world experiments carried out with an ActivMedia Robotics Pioneer 2 robot in populated office environments. The robot was equipped with a SICK LMS laser range finder to perceive its surroundings. The experiments are designed to illustrate how our robot follows people in its surroundings when the environment is crowded.

In the first experiment, the robot moves through a corridor that is empty in the beginning. While the robot is driving, two people walked through as shown in Figure 4. The robot reliably selects people who are suited as a leader and does not consider them as obstacles as conventional approaches would do. It chooses a person to follow as long as someone walks along the corridor towards its goal.

In the second experiment, we test our approach with up to six people walking in a corridor, entering and leaving rooms, or just standing around. The robot is supposed to drive to the other end of the corridor and to return. The people acted naturally as they do in office buildings, railway stations, hospitals, or museums. Figure 5 shows snapshots of this experiment. As before, the robot always selects appropriate people to follow and does not plan trajectories around them which would be a non-humanlike behavior.

The complexity of our planning system corresponds to the one of A* multiplied by the number of iterations. Since the robot can only observe people in its local surrounding, the number of tracked people and therefore the number of iterations is strictly limited. During the experiment, the planning operation required between two and eight iterations and took up to 250 ms on a standard laptop computer. Even with a high number of people in the surrounding of the robot, our system reliably tracked the humans and made use of nearly all possibilities to follow a person.

While testing our system with six subjects, we logged the number of steps needed by our iterative path planning algorithm. We removed all situations in which the robot was able to plan a path directly without following a person. Table I

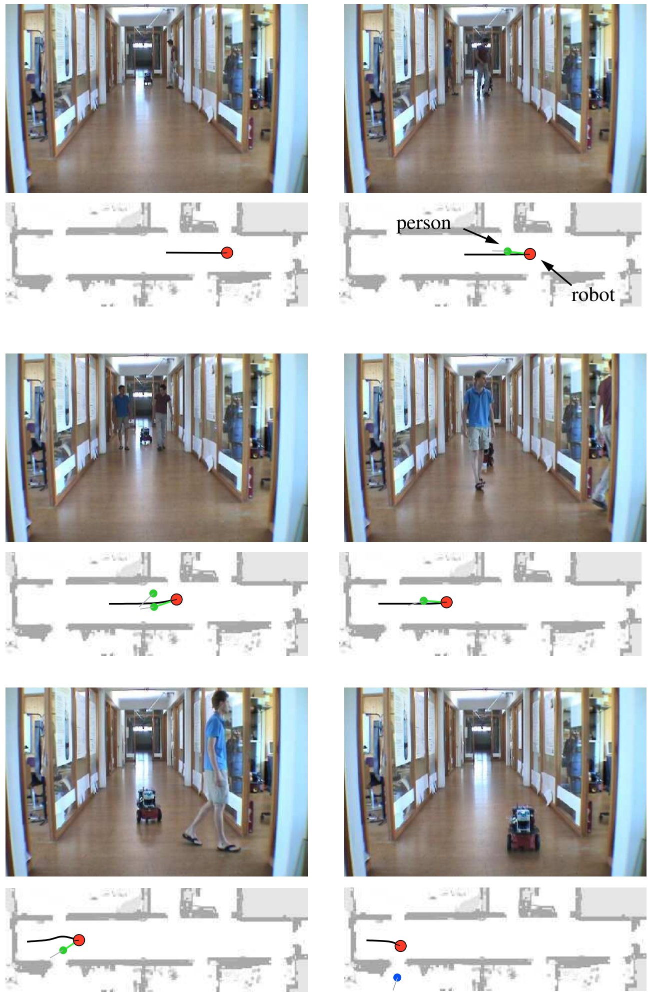

Fig. 4. The robot moved through the corridor and while moving, people entered the corridor blocking its path. Green (light gray) circles indicate people the robot considered to follow and the green line shows its decision. As can be seen, the robot safely navigates through the corridor following people whenever it appears to be appropriate.

TABLE I

NUMBER OF REQUIRED ITERATION STEPS WHILE PLANNING PATHS SURROUNDED BY UP TO 6 PEOPLE.

# iteration	frequency	relative frequency
2	316	81.4%
3	39	10.1%
4	9	2.3%
5	4	1.0%
not terminated	20	5.2%

summarizes the required planning steps. The majority planning operations needed exactly two steps. In around 5% of the cases, our iterative planning approach could not find a solution. This was the case when no plan existed or of the iterative algorithm did not produce an admissible solution.

Additionally we tested to shoo away path blocking people in simulation. We modified our simulator to let a person randomly go away if the robot approaches her or not. Figure 6 shows an example of both cases.

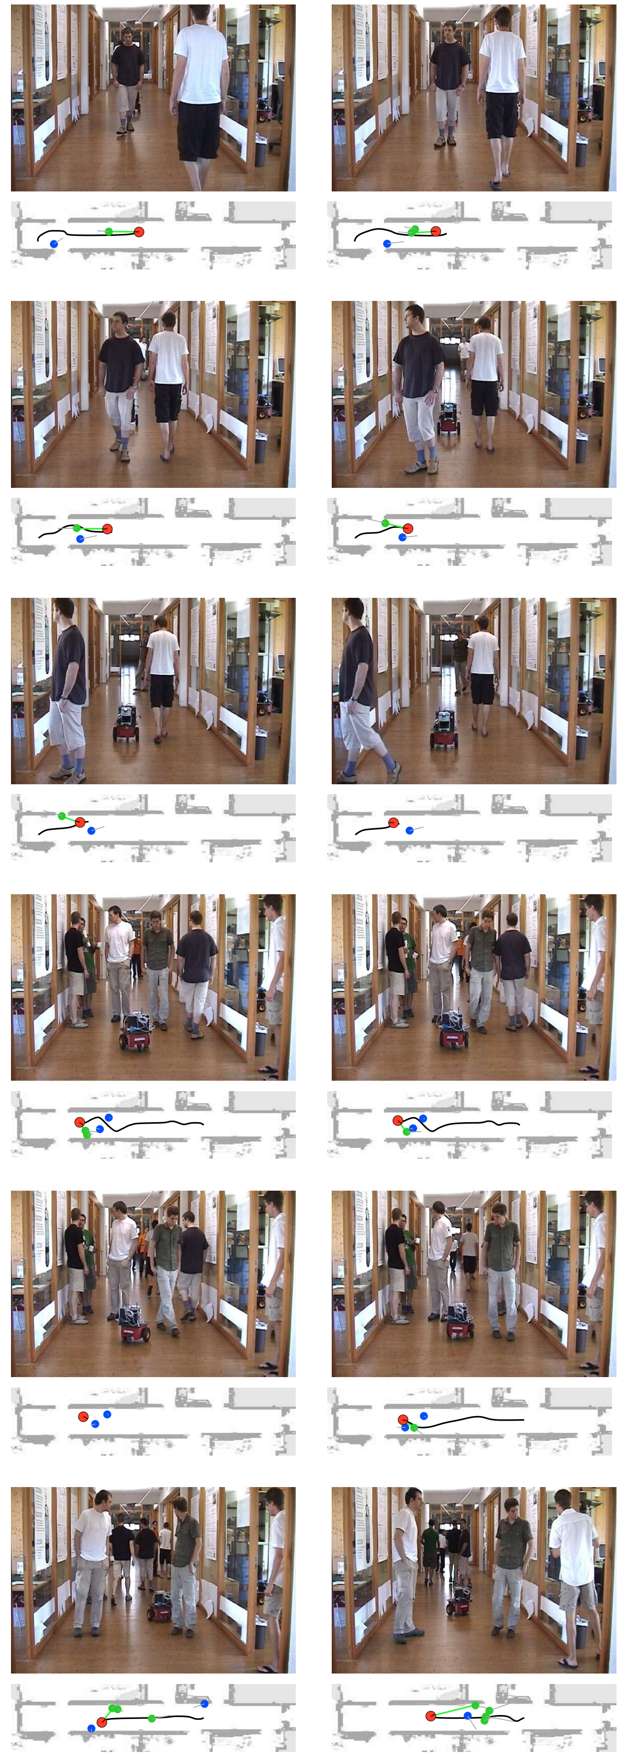

Fig. 5. Experiments with multiple and differently moving people. The robot always selected people that move in the corrected direction – that is in unison with the robot's target location.

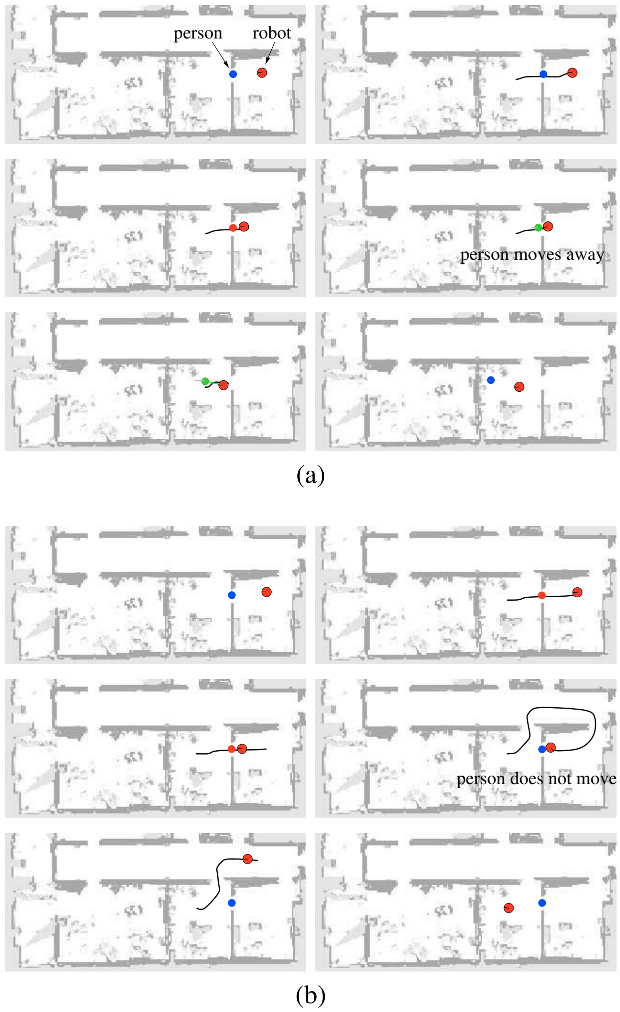

Fig. 6. The robot tries to shoo away a person that is blocking a narrow passage. In situation (a) the person frees the passage so that the robot can move through whereas in (b) the robot has to choose the detour to reach its target location.

VI. Conclusions

In this paper, we presented a novel navigation approach that generates human-like motion behavior for mobile robots in highly populated environments. The approach detects and tracks people in the surroundings of the robot and integrates this knowledge into the planning process. Compared to related work, we do not plan a path around people but try to identify and follow individuals or groups of people that appear to move towards the same goal. This is achieved by an iterative planning and person selection approach based on the A* algorithm that decides which subjects the robot should follow. Our robot furthermore considers actions to shoo people away that block narrow passages towards the desired goal or to avoid significant detours. The approach has been implemented and tested on a real robot in populated environments. The robot shows the ability to move with people in a similar way humans do. Such a technique is advantageous in highly populated environments such as museums, trade fairs, hospitals or public spaces.

Acknowledgment

This work has partly been supported by the DFG under contract number SFB/TR-8, by the EC under contract number FP6-IST-045388-INDIGO and by the German Federal Ministry of Education and Research (BMBF) under grant no. 01IME01F-DESIRE.

References

[1] K.O. Arras, Ó. Martínez Mozos, and W. Burgard. Using boosted features for the detection of people in 2d range data. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2007.

[2] Y. Bar-Shalom and X.-R. Li. *Multitarget-Multisensor Tracking: Principles and Techniques*. YBS Publishing, Storrs, USA, 1995.

[3] M. Bennewitz. *Mobile Robot Navigation in Dynamic Environments*. PhD thesis, University of Freiburg, Dept. of Computer Science, 2004.

[4] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun. Learning motion patterns of people for compliant robot motion. *Int. Journal of Robotics Research*, 24(1), 2005.

[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), 2000.

[6] A. Fod, A. Howard, and M.J. Mataric. Laser-based people tracking. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2002.

[7] A. Foka and P. Trahanias. Predictive autonomous robot navigation. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 490–495, 2002.

[8] G. Grisetti, C. Stachniss, and W. Burgard. Improved techniques for grid mapping with rao-blackwellized particle filters. *IEEE Transactions on Robotics*, 23(1):34–46, 2007.

[9] B. Jensen. *Motion Tracking for Human-Robot Interaction*. PhD thesis, École Polytechnique Fédérale de Lausanne, 2004.

[10] 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. Conf. on Robotics & Automation (ICRA)*, 2001.

[11] J. Mínquez and L. Montano. Nearness diagram navigation (nd): Collision avoidance in troublesome scenarios. *IEEE Transactions on Robotics and Automation*, pages 154–164, 2004.

[12] M. Montemerlo and S. Thrun. Conditional particle filters for simultaneous mobile robot localization and people tracking. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2002.

[13] R. Philippsen, B. Jensen, and R. Siegwart. Toward online probabilistic path replanning in dynamic environments. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2006.

[14] E. Prassler, J. Scholz, and P. Fiorini. A robotic wheelchair for crowded public environment. *IEEE Robotics & Automation Magazine*, 8(1):38–45, 2001.

[15] M. Scheutz, J. McRaven, and G. Cserey. Fast, reliable, adaptive, bimodal people tracking for indoor environments. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2004.

[16] D. Schulz, W. Burgard, D. Fox, and A.B. Cremers. People tracking with a mobile robot using sample-based joint probabilistic data association filters. *Int. Journal of Robotics Research*, 22(2):99–116, 2003.

[17] R. Siegwart, K.O. Arras, S. Bouabdallah, D. Burnier, G. Froidevaux, X. Greppin, B. Jensen, A. Lorotte, L. Mayor, M. Meisser, R. Philippsen, R. Piguet, G. Ramel, G. Terrien, and N. Tomatis. Robox at Expo.02: A large-scale installation of personal robots. *Journal of Robotics & Autonomous Systems*, 42(3-4):203–222, 2003.

[18] C. Stachniss and W. Burgard. An integrated approach to goal-directed obstacle avoidance under dynamic constraints for dynamic environments. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 508–513, 2002.

[19] 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 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pages 518–523, 1995.

[20] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Haehnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. Probabilistic algorithms and the interactive museum tour-guide robot Minerva. *Int. Journal of Robotics Research*, 19(11), 2000.

[21] E.A. Topp and H.I. Christensen. Tracking for following and passing persons. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2005.