Time Dependent Planning on a Layered Social Cost Map for Human-Aware Robot Navigation

Marina Kollmitz*, Kaijen Hsiao†, Johannes Gaa‡ and Wolfram Burgard*
*Department of Computer Science, University of Freiburg, 79110 Freiburg, Germany
{kollmitz,burgard}@informatik.uni-freiburg.de
†Bosch Research and Technology Center (RTC) North America, Palo Alto, CA 94304, USA
kaijen.hsiao@us.bosch.com
‡Institute of Mechatronic Systems, Leibniz Universität Hannover, 30167 Hanover, Germany
johannes.gaa@imes.uni-hannover.de

***Abstract*—As robots make their way into our everyday lives, new behavioral concepts are needed to assure their acceptance as interaction partners. In the presence of humans, robots are required to take safety as well as human comfort into account. This paper presents a novel, planning-based approach for social robot navigation. It uses predicted human trajectories and a social cost function to plan collision-free paths that take human comfort into account. It furthermore employs time dependent, kinodynamic path planning to reason about human motion over time and to account for the kinematic and dynamic constraints of a robot. Our approach generates paths that exhibit properties similar to those used in human-human interaction, such as waiting for a human to pass before continuing along an intended path, avoiding getting too close to another human's personal space, and moving out of the way when blocking a human's path. In extensive experiments carried out with real robots and in simulation we demonstrate the performance of our approach.**

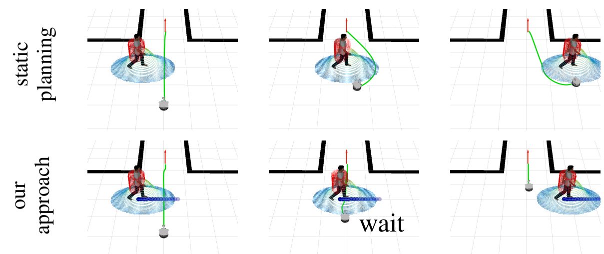

Fig. 1. Comparing static social planning and our approach for a path crossing scenario. With static planning, the robot frequently changes its navigation plans and takes great detours to circumvent the person. In contrast, our approach anticipates the motion of the person and waits in front of the intersection to let him pass.

I. Introduction

Over the past years human-robot interaction has become a field of growing research interest. Robots are increasingly moving away from being single-purpose machines deployed in environments configured especially for them, and towards being intelligent, multi-purpose agents that share their workspace with humans. Taking well-known approaches from the field of autonomous robot navigation can enable robots to navigate in a collision-free manner among people. However, robots need to move beyond simple collision avoidance to ensure that humans do not feel intimidated or annoyed by their presence. Requirements for human comfort and the feeling of safety are particularly pronounced within home environments, for which the presented approach is designed, because robots can directly interfere with a very private and intimate part of a person's life.

Many approaches for human-aware navigation in home environments try to improve human comfort by maintaining an appropriate distance to people. At the same time, many of them do not account for the human motion over time. If people move within the workspace, the robot has to change its navigation behavior frequently due to the false assumption of a static environment as visualized in Fig. 1. As a result, static planning in the presence of humans often produces navigation behavior that lacks consistency and appears confusing and unnatural to people.

The contribution of this paper is a novel approach to human-aware navigation that combines time dependent, search-based path planning with dynamic, social cost maps containing costs based on predicted human trajectories and a Gaussian social cost model. This approach generates paths that exhibit properties similar to those used in human-human interaction, such as waiting for a human to pass before continuing along an intended path, avoiding getting too close to another human's personal space, and moving out of the way when blocking a human's path. Our code is available at `https://github.com/marinaKollmitz/human_aware_navigation`.

II. Related Work

Research in the context of human-human interaction indicates that appropriate interaction distances are important to ensure human comfort [6]. Various studies reveal that the preferred interaction distances are generally also valid for human-robot interaction scenarios [9, 20] and stress the importance of the spatial relationship of robots towards humans for their acceptance as interaction partners. Consistent and predictable robot behavior, characterized by velocity adaptation and waiting behavior (slowing or even pausing to avoid crossing a human's path) rather than by performing large path deviations, is another important aspect of social robot navigation [11, 14]. Furthermore, people expect robots in the home to show considerate and polite behavior and to take on a rather devoted role [2].

Various approaches for human-aware navigation use social cost functions or potential fields to account for social constraints on the (global) path planning level. Kirby et al. [10] use a combination of cost functions to represent social constraints such as the social space preferences of people and a tendency to pass a person on the right side of a corridor. Sisbot et al. [18] include aspects of safety, comfort, visibility and hidden zones into their cost model. A major drawback of the above mentioned approaches is that they do not account for the motion of humans over time. By using static planning, the approaches treat people in the workspace as static instead of dynamic obstacles. The robot has to change its navigation behavior frequently to adapt to the unforeseen changes in the environment. As a result, static planning often produces inconsistent robot behavior that might feel unnatural and confusing for people.

Global planning approaches that account for the motion of dynamic obstacles (e.g. [8, 13, 16]) are generally insufficient to produce socially acceptable navigation behavior, because they do not take social constraints into account. Göller et al. [5] use A* planning on an (x, y, t) occupancy grid in combination with reactive local planning to plan safe robot paths inside a populated supermarket environment. Bennewitz et al. [1] use human prediction hypotheses based on learned motion patterns from observed trajectories and planning in the time-configuration space (x, y, t) to produce socially acceptable navigation behavior. Both approaches incorporate the possibility of velocity adaptation and waiting to solve situations of conflict. However, they do not represent the proxemic preferences of people and might therefore produce navigation paths that feel offensive and uncomfortable to people. Furthermore, the approaches do not account for kinematic and dynamic constraints of robots and are therefore only suitable for holonomic robots with negligible acceleration limits.

Scandolo and Fraichard [17] and Svenstrup et al. [19] use time dependent social constraints and randomized kinodynamic planning to account for the motion of people and their social space preferences over time. While randomized planning is generally less computationally expensive than deterministic planning as employed in our approach, no guarantee to optimality can be given. Since we use a cost function which weights different path aspects to adjust the navigation behavior, we found it less appropriate to use a planner that generally does not optimize the cost function.

Learning-based approaches try to reproduce appropriate navigation behavior by observing humans. Kuderer et al. [12] use features to capture important properties of human navigation behavior and estimate the feature weights by machine learning. The resulting joint collision avoidance behavior is socially acceptable and even unavoidable for navigation among crowds of people in public spaces. However, we aim for a more conservative and obedient robot behavior for navigation within the home.

Finally, some navigation approaches account for obstacles on the local planning level, e.g. the Velocity Obstacle approach [3] or the Dynamic Window Approach (DWA) [4]. Some approaches are specifically designed for navigation among humans, such as the Social Force Model [7]. Weinrich et al. [21] use belief distribution maps composed by a life-long learning technique to predict human motion and use the spatio-temporal occupation probabilities to formulate an additional objective inside the DWA. However, local planning approaches are generally not able to solve complex interaction situations in a consistent way due to the limited look-ahead.

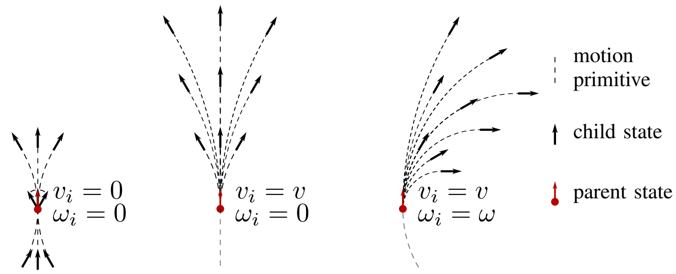

Fig. 2. The state transitions of the search graph are directly executable motion primitives. They are constructed according to the kinematic model of the robot and satisfy its dynamic constraints.

III. Human-Aware Path Planning

Our human-aware navigation approach aims to improve human comfort during navigation in the presence of humans, focusing on home environments that typically contain a limited number of humans but may be narrow or cluttered. We attempt to satisfy the social space preferences during planning by introducing additional social constraints to the path planning process. The constraints are supposed to prevent the robot from operating in the intimate zone of people and prefer the personal zone for interaction between $0.45\,\mathrm{m}$ and $1.2\,\mathrm{m}$ from the person's center, as suggested by HRI related research [9, 20]. Furthermore, a main focus lies on consistent robot paths to improve the legibility of the navigation behavior. We concentrate on differential drive motion because holonomic motion contrary to the heading direction might not seem natural and goal directed for people. In general, we aim for a polite and obedient robot navigation behavior that gives priority to humans and includes making way for them.

We use time dependent, deterministic planning (A*) to reason about the spatial relationship of a robot and humans with respect to time. Section III-A describes the search graph composition in more detail. A social cost model described in Section III-B defines additional social constraints for planning. We use a layered cost map to efficiently combine the time dependent social constraints and static environment constraints as described in Section III-C. Our algorithm optimizes the path in terms of social comfort, path length, execution time and environment constraints as described in Section III-D. A heuristic (Section III-E) accelerates the search by estimating the remaining costs to reach the goal. Finally, a planning timeout described in Section III-F ensures a constant planning frequency.

A. Search Graph

Adding time to the planning space introduces additional complexity to the planning process. The planned, time dependent trajectories have to be executable to ensure that the robot can follow the planned trajectory in time. The reachable robot configurations within a time interval Δt depend on the robot's current pose (x, y, φ), its velocity (v, ω) and its kinematic and dynamic constraints. To produce feasible, time dependent paths for differential drive robots with inertial constraints we

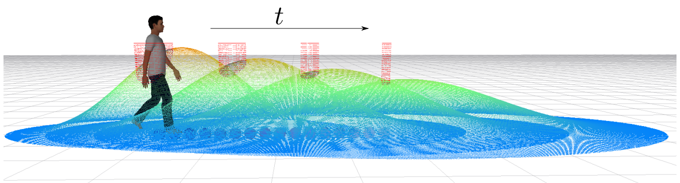

Fig. 3. Time dependent decay functions overlay the Gaussian-based social cost model to account for prediction uncertainties.

plan with executable state transitions inside a state space of $C = (x, y, \varphi, v, \omega, t)$.

The search graph is constructed according to a discrete-time model with a constant time interval Δt for each expansion and a finite set of directly executable actions u_i. The actions u_i in the action set

$$U = \{(e_x a_x, e_\phi a_\phi) | e_x, e_\phi \in \{-1, 0, 1\}\} \tag{1}$$

are a combination of forward and angular accelerations a_x and a_φ. The resulting motion primitives ρ_i each represent constant acceleration, constant deceleration or constant velocity, respectively, for both velocity components v and ω as shown in Fig. 2. Expanding the graph with motion primitives produces a reachability tree that is forced into state lattice representation [15] by regular cell decomposition to enable the pruning of states that are similar to previously expanded states.

B. Social Cost Model

Our approach uses a social cost model to represent the social space preferences of humans as visualized in Fig. 3. Like in previous approaches [10, 18], it is based on a Gaussian distribution with an amplitude A and different standard deviations σ_x and σ_y along a person's front and side. The Gaussian distribution is displaced from the person's center by Δx and Δy to grant special consideration to the sensitive area directly in front of a person and to reflect social preferences like passing on the right side of a person. Furthermore, a non-traversable area within a forbidden radius r_0 around the person's center (visualized as a red pillar in Fig. 3) accounts for the physical outline of a person. Paths traversing this area are pruned directly during planning because they would result in collision. We use decay factors for the amplitude d_A, the variances d_σ and the forbidden obstacle radius $d_{r,0}$ to account for prediction uncertainties. The decay factors induce a linear decrease of the amplitude and forbidden obstacle radius and a linear increase of the variances of the social cost function over time. They enlarge the area that is potentially affected by the presence of people over time, but lower the costs of occupying the influenced space to avoid that the robot becomes overcautious.

C. Dynamic Cost Map

We use a layered dynamic cost map to represent static environment constraints and time dependent social constraints. It consists of one static layer and multiple dynamic layers as visualized in Fig. 4. All layers are two-dimensional grids that define navigation constraints as costs for occupying or traversing space. The static layer contains costs due to environment

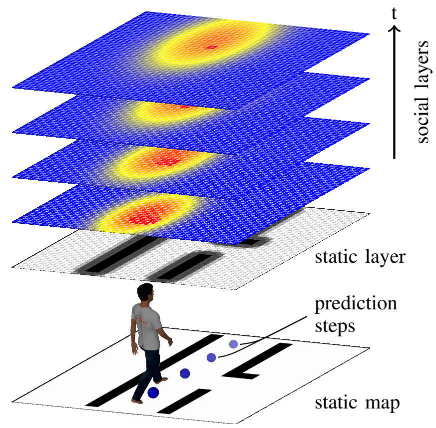

Fig. 4. The dynamic cost map consists of one static layer that represents environment constraints and multiple social layers. Each social layer represents the social navigation constraints for one prediction time step.

constraints obtained from the static map, sensor observations, etc.. The dynamic layers incorporate costs related to social constraints according to the social cost model presented above. Each dynamic layer represents one snapshot of the predicted human trajectory from time step i until step $(i + 1)$. During planning, the costs of occupying or traversing space at time t_i are obtained from the dynamic cost map as a combination of the costs from the static layer and the i-th dynamic layer. If the queried time t_i is outside the time frame represented by the dynamic layers, the costs from the static map are returned. The represented time frame should be as large as possible to capture complex interaction scenarios. However, a large prediction look-ahead increases the computational effort during planning.

The dynamic cost map allows for an efficient representation of time dependent navigation efforts since the query of costs is simplified to a lookup operation in a constant array. However, large memory requirements arise from large maps, fine space or time resolutions or a large prediction look-ahead. The cost map representation separates the planning procedure and the definition of planning constraints. Therefore, it is easy to exchange or enhance our cost model or incorporate a different prediction technique if desired by populating the cost map accordingly.

D. Optimization

The A* search algorithm optimizes the costs along the planned path according to a cost function over navigation actions. Our algorithm evaluates the planned paths in terms of path execution time, path distance, static environment constraints and social constraints. Each of these factors is represented by a cost c_i. The cost for the path execution time is proportional to the required time steps to reach the goal. The path length is evaluated based on the linear and diagonal transitions in the cost map grid required to follow a planned path[1]. The static environment costs as well as the

[1]Our heuristic estimates the remaining costs to the goal by an eight-connected expansion inside the cost map grid (see Section III-E). The path length is also evaluated according to grid transitions for consistency.

cost function weights	
path execution time	400/s
path length	1600/m
static env. max value	255
social cost max value	255

social cost model	
amplitude A	255
standard deviation σ_x	0.25 m
standard deviation σ_y	0.2 m
forbidden radius r_0	0.25 m
offset Δx	0.2 m
offset Δy	0.0 m
decay factor d_a	-0.1/s
decay factor d_σ	0.05/s
decay factor $d_{r,0}$	-0.1/s

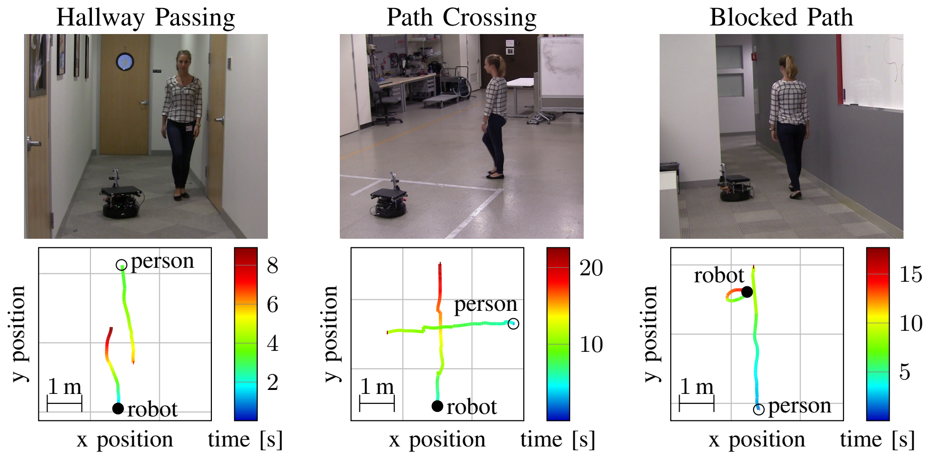

Fig. 5. Navigation experiments. Left: cost function parameters. Right: Human and robot trajectories for three standard navigation scenarios. The trajectories are colored to visualize the time dimension. Note that path jumps and irregularities are caused by the localization and the people detection module.

costs of occupying socially sensitive space are obtained from the dynamic cost map presented above, according to the grid cells the path traverses. The resulting cost function

$$\mathcal{C} = \sum w_i \cdot c_i \qquad (2)$$

is a weighted sum of the individual cost components c_i. The weighting factors w_i specify the priorities of the related path properties. Path length and path execution time weighting factors rate the priority of velocity adaptation and waiting behavior against path deviations in the presence of humans. The weighting factor for social costs determines the tradeoff between necessary detour or waiting time and closer proximity to people.

E. A Heuristic*

The remaining costs to reach the goal are estimated based on the simplifications that no people are present in the environment, as well as that the robot is holonomic and does not have inertial constraints. A planar, eight-connected Dijkstra expansion is conducted on the static layer of the dynamic cost map to find the shortest possible path to the goal and the minimal static environment costs in the absence of people. The time to reach the goal is estimated by assuming that the robot can execute the planned path at maximum velocity. The heuristic is admissible and consistent and therefore satisfies the optimality criterion.

F. Planning Timeout

We use a planning timeout to limit the time allocated for planning and ensure a constant planning frequency. Once the planning time runs out, the algorithm returns the path to the best expanded state, which ideally is the goal state. If the goal state was not reached during search, it returns the lowest cost path that reached the goal at any orientation and velocity. In case no state reached the goal position, the algorithm chooses the path to the state that would be expanded next during search. Since A* always expands the state that produces the lowest expected costs to reach the goal, the selected path is the most promising to advance to the goal. Kushleyev and Likhachev [13] propose to dilute the planning space and plan statically once the states exceed a certain time bound. Instead of limiting the time horizon within the planning space, we limit the time allocated for planning. Thereby, we can ensure a constant planning frequency but produce complete feasible navigation paths if the associated planning time allows.

IV. Experiment Results

We conducted extensive experiments to evaluate the performance of our human-aware navigation approach. Section IV-A presents the performance of our navigation system for three standard navigation scenarios with one person. We compare our approach to a static, social cost map based planning approach in Section IV-B.

A. Real Robot Experiments

We conducted three experiments with the Turtlebot2 platform to evaluate the performance of the implemented human-aware navigation approach for typical navigation scenarios. A Kalman filter based detection module estimated the position and velocity of the interacting pedestrian from scans of the on-board 2D laser range finder. The future trajectories of the person were estimated with a constant velocity model. Our approach used a map resolution of 0.05 m, a time resolution of 0.5 s and 15 prediction steps (7.5 s look-ahead). The planning frequency was kept constant at 2 Hz. The parameters of the social cost function and the cost function weights as well as the resulting robot and human trajectories are presented in Fig. 5. The cost parameters were chosen to assign a high priority to respecting the person's social space, and to prefer direct paths with velocity adaptations instead of taking great path deviations.

The first experiment tested a hallway passing situation. The robot moves to the left hand side of the corridor to avoid the approaching person. It keeps a sufficient lateral passing distance to respect the spatial preferences of the human.

The second experiment tested the navigation performance of our approach for a path crossing situation. The robot proceeds to the intersection and waits for the person to pass before proceeding to the goal[2]. It chooses a direct path and waits for the human to pass, as expected of robots in path-crossing situations [14], instead of taking a great detour.

The third experiment evaluated the navigation behavior for a situation where the robot initially rests at a conflicting area. As the person approaches the robot, it moves to the side in order to let the person pass. Once the person has left the area of conflict, the robot returns to its resting position. The presented behavior is achieved by requiring a valid path for at least the predicted trajectory window, not only until the goal is reached. Thus, the robot can anticipate and avoid future situations of conflict that occur after reaching the goal.

B. Static Planning Comparison

Two sets of experiments evaluate the performance of our approach against 2D static deterministic planning[3] in populated environments. The experiments were conducted in simulation to ensure equal conditions for both approaches. Furthermore, we assumed that the robot always knows the position and velocity of people in the environment. The maximum robot velocity was set to $0.4\,\mathrm{m/s}$ for all runs and both approaches used the social cost model presented in Section III-B. The static planning approach produced paths at a frequency of $2\,\mathrm{Hz}$. The planning time was not limited for our approach to examine how long it takes to solve the complete navigation situation. We enlarged the planning look-ahead to $30\,\mathrm{s}$ and removed the decay factors of the social cost model to maintain the correct representation of social constraints over time without replanning. Otherwise, the planner configuration was kept like in the real robot experiments (Section IV-A).

The first set of experiments tested a crossing situation with one person as depicted in Fig. 6(a). The set contains 50 runs at varying person velocities that were generated randomly from a Gaussian distribution with $0.4\,\mathrm{m/s}$ mean and $0.1\,\mathrm{m/s}$ standard deviation. The resulting robot trajectories and a comparison of characteristic path properties is presented in Fig. 7. The behavior of the standard navigation approach is characterized by large, inefficient path deviations to circumvent the person. Furthermore, the robot collided twice with a person while trying to cross in front. In contrast, slowing down or waiting in front of the intersection to let the person pass are the predominant navigation decisions in our approach. The robot kept an appropriate distance to the person during all runs and on average executed shorter paths in less time. Our approach took on average $0.8\,\mathrm{s}$ to generate a path ($1.5\,\mathrm{s}$ standard deviation).

The second set of experiments, visualized in Fig. 6(b), tested a hallway navigation scenario with two people. Person 1 started in front of the robot and moved along the hallway at a velocity that was generated randomly from a Gaussian distribution with $0.2\,\mathrm{m/s}$ mean and $0.05\,\mathrm{m/s}$ standard deviation. Person 2 moved at a velocity of $0.4\,\mathrm{m/s}$ towards person 1 from the opposite side of the hallway. The results for the hallway experiment are summarized in Fig. 7. With the standard navigation approach, the robot directly tried to overtake person 1 to proceed to the goal without waiting for person 2 to pass.

[2]The waiting position is determined by the order at which the states are expanded during A* planning, since waiting at any point in front of the intersection will produce the same costs.

[3]ROS indigo navigation stack: `http://wiki.ros.org/navigation`

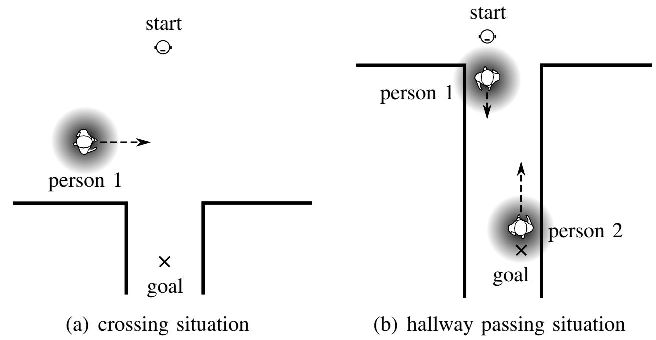

Fig. 6. Experiment setup for the comparison experiments.

As person 2 moved along, the robot tried to pursue a path between both people. However, it was not able to react in time and clear the area of conflict in a socially compliant way. The robot collided with person 2 in 25 runs and was generally not able to respect the intimate space of person 2. In a real world scenario, the robot would probably have disturbed both people and forced them to deviate from their intended paths to avoid collisions and unsocial behavior. In contrast, our approach made the robot slow down or wait to keep behind person 1 until person 2 had passed. Only after person 2 cleared the area of conflict did the robot proceed to the goal by overtaking person 1. The robot was able to keep an appropriate distance to both people at all runs and reached the goal after a comparable execution time and path length. Our approach took on average $3.5\,\mathrm{s}$ to generate a path ($2.0\,\mathrm{s}$ standard deviation).

V. Conclusion

In this paper, we presented a planning-based approach to human-aware navigation that uses a social cost function and planning in time to improve mobile robot navigation behavior in populated environments. In extensive experiments we demonstrated that our approach can reliably generate consistent paths that respect the social space preferences of one or more people in the environment. The experiments revealed advantages of our approach compared to static planning approaches. To handle the increased computational effort compared to planning in the 2D space, we introduced a planning timeout to ensure a constant planning frequency. In future work, we will consider faster replanning algorithms or a reduction of the planning dimension for robots with negligible dynamic constraints to further reduce the planning time.

References

[1] M. Bennewitz, W. Burgard, and S. Thrun. Adapting navigation strategies using motions patterns of people. In *Proc. of the IEEE Int. Conf. of Robotics and Automation (ICRA)*, 2003.

[2] K. Dautenhahn, S. Woods, C. Kaouri, M. L. Walters, K. L. Koay, and I. Werry. What is a robot companion - friend, assistant or butler? In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2005.

[3] P. Fiorini and Z. Shiller. Motion planning in dynamic environments using velocity obstacles. *Int. Journal of Robotics Research*, 17(7), 1998.

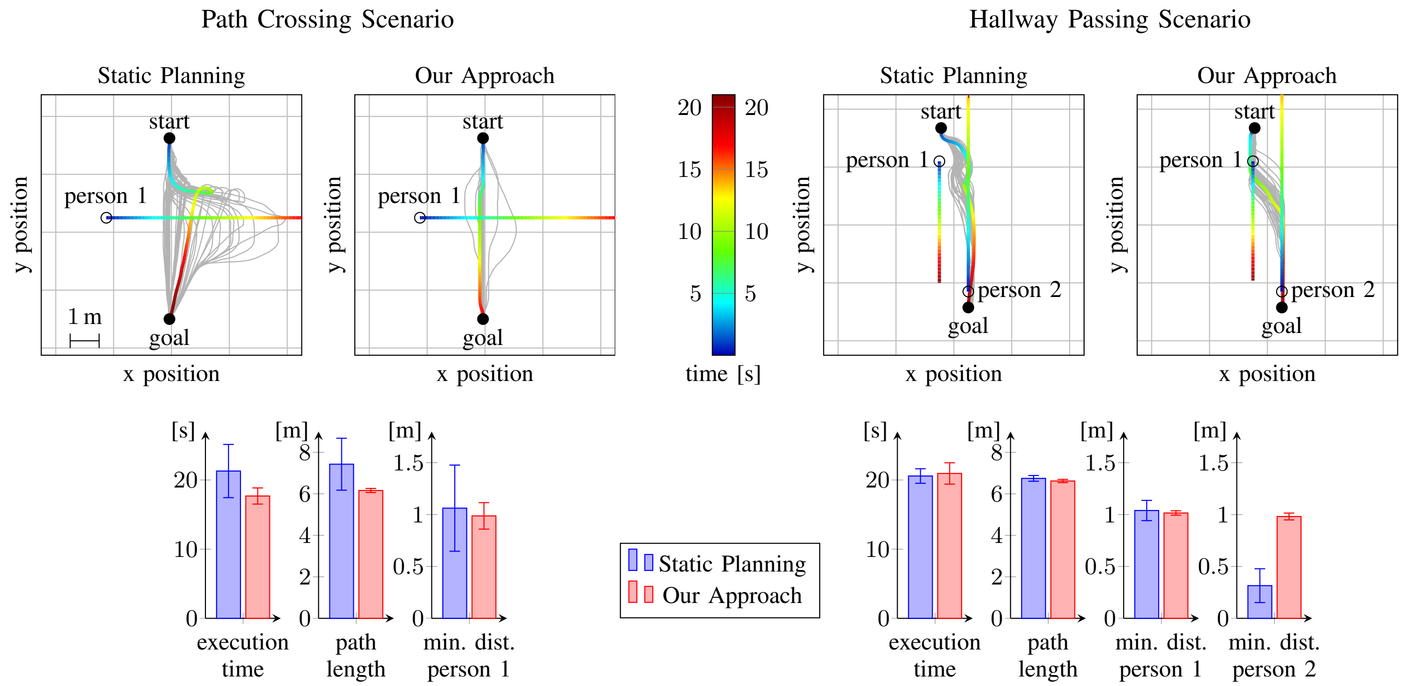

Fig. 7. Comparison of a static planning approach and our approach for the crossing (left column) and hallway passing (right column) scenario. The top charts visualize 50 runs at varying person velocities. One run at the mean person velocity is highlighted and colored to visualize the time dimension. The bar charts compare the average values of characteristic path properties for both approaches, the standard deviation is visualized by error bars.

[4] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine*, 4(1), 1997.

[5] M. Göller, F. Steinhardt, T. Kerscher, J. M. Zöllner, and R. Dillmann. Proactive avoidance of moving obstacles for a service robot utilizing a behavior-based control. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2010.

[6] E. T. Hall. *The hidden dimension*. Doubleday, 1966.

[7] D. Helbing and P. Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51, 1995.

[8] D. Hsu, R. Kindel, J.-C. Latombe, and S. Rock. Randomized kinodynamic motion planning with moving obstacles. *Int. Journal of Robotics Research*, 21(3), 2002.

[9] H. Huettenrauch, K. Severinson Eklundh, A. Green, and E. A. Topp. Investigating spatial relationships in human-robot interaction. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2006.

[10] R. Kirby, R. Simmons, and J. Forlizzi. Companion: A constraint optimizing method for person-acceptable navigation. In *Proc. of the IEEE Int. Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2009.

[11] T. Kruse, P. Basili, S. Glasauer, and A. Kirsch. Legible robot navigation in the proximity of moving humans. In *Proc. of the IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*, 2012.

[12] M. Kuderer, H. Kretzschmar, C. Sprunk, and W. Burgard. Feature-based prediction of trajectories for socially compliant navigation. In *Proc. of Robotics: Science and Systems (RSS)*, 2012.

[13] A. Kushleyev and M. Likhachev. Time-bounded lattice for efficient planning in dynamic environments. In *Proc. of the IEEE Int. Conf. of Robotics and Automation (ICRA)*, 2009.

[14] C. Lichtenthäler, A. Peters, S. Griffiths, and A. Kirsch. Social navigation-identifying robot navigation patterns in a path crossing scenario. In *Social Robotics*. Springer, 2013.

[15] S. Pancanti, L. Pallottino, D. Salvadorini, and A. Bicchi. Motion planning through symbols and lattices. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2004.

[16] S. Petti and T. Fraichard. Safe motion planning in dynamic environments. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2005.

[17] L. Scandolo and T. Fraichard. An anthropomorphic navigation scheme for dynamic scenarios. In *Proc. of the IEEE int. Conf. on Robotics and Automation (ICRA)*, 2011.

[18] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Siméon. A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, 23(5), 2007.

[19] M. Svenstrup, T. Bak, and H. Andersen. Trajectory planning for robots in dynamic human environments. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2010.

[20] M. L. Walters, K. Dautenhahn, K. L. Koay, C. Kaouri, R. te Boekhorst, C. Nehaniv, I. Werry, and D. Lee. Close encounters: spatial distances between people and a robot of mechanistic appearance. In *Proc. of the IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids)*, 2005.

[21] C. Weinrich, M. Volkhardt, E. Einhorn, and H.-M. Gross. Prediction of human collision avoidance behavior by lifelong learning for socially compliant robot navigation. In *Proc. of the IEEE int. Conf. on Robotics and Automation (ICRA)*, 2013.