

Socially-Aware Robot Navigation: A Learning Approach

Matthias Luber Luciano Spinello Jens Silva Kai O. Arras

Abstract—The ability to act in a socially-aware way is a key skill for robots that share a space with humans. In this paper we address the problem of socially-aware navigation among people that meets objective criteria such as travel time or path length as well as subjective criteria such as social comfort. Opposed to model-based approaches typically taken in related work, we pose the problem as an unsupervised learning problem. We learn a set of dynamic motion prototypes from observations of relative motion behavior of humans found in publicly available surveillance data sets. The learned motion prototypes are then used to compute dynamic cost maps for path planning using an any-angle A* algorithm. In the evaluation we demonstrate that the learned behaviors are better in reproducing human relative motion in both criteria than a Proxemics-based baseline method.

I. INTRODUCTION

With robots entering our daily lives, their ability to perceive, understand, and act in a socially conform way is becoming a key requirement for many application scenarios. Research in the area of socially-aware navigation and manipulation is typically taking a model-based approach, either with manually designed models or models from social psychology and cognitive science [13, 16, 14, 9, 12, 17]. Although this appears to be a compelling approach, the use of such models in robotics is potentially problematic due to an inherent *methodological gap*: the experimental method of these fields isolates and studies an aspect of interest under controlled conditions. In contrast, the robotics community seeks to take robots out of the controlled laboratory setting into the complex real world full of unpredictable events.

Examples include the work of Pacchierotti *et al.* [13] on a navigation system for hallways based on Hall's Proxemics theory [4]. However, Proxemics is the study of measurable distances between people as they *interact*, not walk past someone. The authors in [16] also use Proxemics in combination with a back space model that describes the discomfort of trajectories that pass behind a person for a dynamically navigating robot in a populated environment. However, the latter model was derived from behavior of male subjects in a lavatory context [11] and it is questionable if it extends to robots in a general crowd. A robot collision avoidance system based on a social force model [6] is proposed in [17].

All authors are with the Social Robotics Lab, Department of Computer Science, University of Freiburg, Germany {luber,spinello,silva,arras}@informatik.uni-freiburg.de.

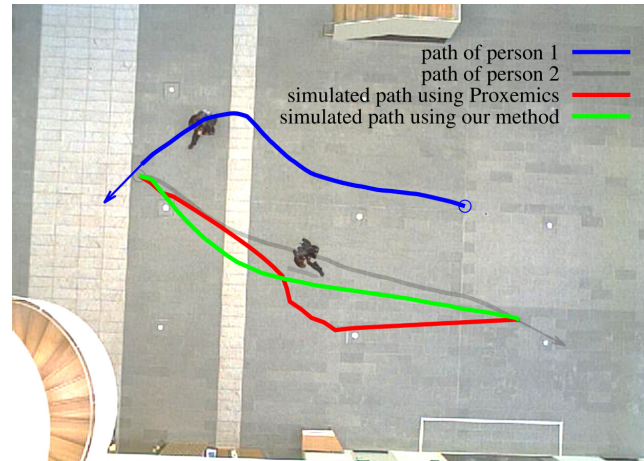


Fig. 1. Socially-aware navigation among humans. Two pedestrians (blue and gray) avoid each other. The learned relative motion prototype (in green) resembles the path of a real person (in gray) while a Proxemics-based baseline path (in red) contains unnecessary heading changes and is longer.

But the model was developed to describe the behavior of *dense* pedestrians crowds.

Other works rely on manually designed ad-hoc models not directly derived from other cross-disciplinary fields. In [12], a standing in line behavior of a robot was engineered. In [9] the authors define cost functions for human preferences, visual constraints and a comfort score to achieve socially acceptable hand-over motions.

In all these examples, it is questionable if the employed models scale to the operational conditions of a robot in the complex real world. We thus believe that socially-aware behavior should be *learned* from real data. We extend the current state of the art by learning a socially-aware navigation model from observations of humans, acquired and used under the same real-world conditions.

The work by Henry *et al.* [7] also takes a learning approach to a similar problem. However, the authors address the problem of learning a planning strategy through streams of pedestrians rather than learning continuous sensory-motor motions. Using inverse reinforcement learning with features such as density and flow direction, their agent learns to plan minimum hindrance paths through simulated pedestrian streams. Here, in contrast, we are addressing the (lower-level) problem of learning continuous motion prototypes that are considered socially acceptable while being task-efficient in terms of time or path length. Most importantly, we learn




Fig. 2. Top view of the environment monitored in the EIPD data set. The size of the area is $15.8m \times 11.8m$. The white arrows depict entry or exit points for people in the scene. The overlaid colored plots show example walking paths of people in the environment.

this from real data on human motion as opposed to simulated pedestrians.

To this end, we take annotated surveillance data sets collected from overhead cameras (see Fig. 2), extract the pedestrian paths, and transform them into a 3D representation. We then learn human relative motion behavior in an unsupervised way which will result in a set of dynamic motion prototypes that enable the robot to compute a dynamic cost map and plan socially acceptable avoidance maneuvers among humans that reflect the learned relative motion behaviors. Finally, we evaluate the models with respect to the human behavior to see how well they generalize to never seen instances of human behavior.

The paper is organized as follows: the next section presents the unsupervised learning approach of relative motion prototypes followed by section III that describes how a dynamic cost map is derived from the learned behavior and how planning is made on such maps. Section IV contains the evaluation of the learned motion prototypes and section V concludes the paper.

II. LEARNING RELATIVE MOTION PROTOTYPES

In this section we present the theory for learning socially-aware relative motion prototypes (RMP). First we give their formal definition thereafter the procedure for unsupervised learning is explained.

Human spatial motion is influenced by many factors such as goal position, obstacles in the way, other people and social norms. The dependencies among these variables are complex in particular when multiple people are involved. Thus, we break down the problem into pairwise relationships and consider the movements of only two persons at once in relation to their goals. The practice of analyzing only pairwise relations between variables is common to ease the modeling of complex many-to-many




Fig. 3. The key parameter used in this work to learn and reproduce socially-enabled behaviours is the angle of approach α_{ij} that is measured between the space-time trajectories of the two persons.

relationships. Examples include graphical models [5, 8].

A relative motion prototype $\mathcal{R}_{i,j}$ describes a relative motion between person π_i and person π_j . Given the two observation sequences \mathbf{z}_i and \mathbf{z}_j of their (x,y) -positions over time t , we define

$$\begin{aligned} d_{i,j}(t) &= \|\mathbf{z}_i(t) - \mathbf{z}_j(t)\| \\ \mathcal{R}_{i,j} &= [d_{i,j}(t_s), \dots, d_{i,j}(t_e)] \end{aligned} \quad (1)$$

with $[t_s, t_e]$ being the duration of the prototype. As we consider sequences of discrete observations, the function R becomes a set of distances

$$\mathcal{R}_{i,j} = \{d_{i,j}(t_s), \dots, d_{i,j}(t_e)\} \quad (2)$$

defined on the interval $\{t_s, \dots, t_e\}$ of discrete time indices. The definition of $\mathcal{R}_{i,j}$ encodes the relative motion between two persons invariant to their absolute position and orientation in the environment.

A. Clustering Sequences of Relative Motion

To learn relative motion prototypes in an unsupervised manner, we take an hierarchical clustering approach. As relative movement sequences can differ in duration and relative speed, we have to define an appropriate distance function able to group similar motion behaviors into the same cluster (further referred to as prototype). Concretely, we use hierarchical agglomerative clustering with average linkage (see [3]) and distance function calculated with a modified Dynamic Time Warping algorithm (DTW, see [15]). DTW is an algorithm originally developed in the field of speech processing for matching sequences of spoken language at different speeds. We extend the regular DTW algorithm by introducing an asymmetry into the weights to each operation. This leads to the asymmetric DTW (aDTW) shown in Algorithm 1. Three operations are considered when two sequences are compared: insertion, deletion and matching. The asymmetric DTW algorithm penalizes insertion or deletion operations depending on the length of the two sequences to compare. If the first sequence is longer compared to the

Algorithm 1: Asymmetric Dynamic Time Warping

```

input : Sequence  $s$  of length  $n$ , sequence  $t$  of
        length  $m$ , penalty  $\beta$ 
output: Distance between sequences  $s$  and  $t$ 
// initialization
for  $i := 1$  to  $n$  do
   $aDTW[i, 0] := \infty$ ;
for  $j := 1$  to  $m$  do
   $aDTW[0, j] := \infty$ ;
 $aDTW[0, 0] := 0$ ;
// compute penalty for operations
if  $n > m$  then
   $c_1 := \beta$ ;  $c_2 := 1$ 
else
   $c_1 := 1$ ;  $c_2 := \beta$ 
// compute similarity
for  $i := 1$  to  $n$  do
  for  $j := 1$  to  $m$  do
     $cost := dist(s[i], t[j])$ ;
     $aDTW[i, j] := cost + \min($ 
       $c_1 \cdot aDTW[i-1, j],$  // insertion
       $c_2 \cdot aDTW[i, j-1],$  // deletion
       $aDTW[i-1, j-1]$  // matching
     $);$ 
return  $aDTW[n, m]$ ;

```

second one insertions are penalized, otherwise it penalizes deletion operations. The insight is that aDTW leads to low costs between two sequences when one of them can be transformed into the other by a simple uniform time rescaling. Moreover, aDTW allows to tune the desired uniformity by adjusting the stiffness parameter β . As opposed to regular DTW, the proposed distance is ideal for clustering sequences in which preservation of the qualitative shape up to time scaling is important. During clustering, we compute the centroids by averaging the two closest relative movement sequences. This is done along the correspondences found by the aDTW procedure via linear interpolation and averaging. The resulting sequence is stretched to the length of the longer sequence.

B. Social Context Model Selection

If all observed motion sequences were clustered at this point, we would obtain meaningless results as the goal and other variables that influence human motion are not considered yet. Thus, in this section, we present a method to account for these factors by modeling and recognizing the social context that gives rise to a certain normative motion behavior.

The problem can be posed as a model selection problem: how many different social context situations need to be distinguished to optimize the trade-off between




Fig. 5. Visualization of a set of learned relative motion prototypes (RMP) encoding five different social context situations. The dashed horizontal black lines show the boundaries of the Personal space as defined in the Proxemics theory. The learned models clearly intrude the Personal space indicating that the Proxemics theory is invalid for walking people. This finding reinforces our motivation to take a learning approach to the problem.

data likelihood and model complexity? If we increase model complexity and consider many situations, the learned prototypes will overly explain the relative motion behavior in each context due to overfitting. If we decrease model complexity and distinguish only few situations, the prototypes generalize well but are poor explanations of the data. To resolve this trade-off, we consider a measure similar to the Bayes Information Criterion (BIC) that selects among a finite set of models to balance the generalization error and model complexity.

Concretely, we use the angle of approach $\alpha_{i,j}$ between the two subjects π_i and π_j as the criterion to quantify and distinguish what we define here to be a social context (see Fig. 3). This choice is motivated by insights from social psychology [1, 11] where indication was found for this to be an important factor in modeling interpersonal comfort and motion behavior. We estimate the angle of approach $\alpha_{i,j}$ by tracking the persons and predicting their future trajectories. To this end, we employ a Kalman filter with state components position, orientation, and velocity and a constant velocity motion model. As soon as the filter is initialized and has reached steady-state (usually after three or four steps), we obtain a stable estimate for $\alpha_{i,j}$. Since we are interested in behavioral changes only caused by nearby subjects, we also compute the closest distance between the extrapolated trajectories and consider only pairs of pedestrians that approach each other closer than a threshold distance.

We then determine the optimal number n of social contexts, that is the optimal number of $\alpha_{i,j}$ -intervals. We start with $n = 1$ context and group all motion sequences into a single RMP generated with all angles of approach $\alpha_{i,j} \in A$ where $A = [0, 2\pi)$. Next, the goodness-of-fit term in the BIC is computed. We continue by increasing n , splitting A into two intervals. The procedure stops when the optimum BIC value is found. We finally define $\mathcal{P} = \{\mathcal{R}_{i,j}^{\alpha_1}, \dots, \mathcal{R}_{i,j}^{\alpha_n}\}$ to be the set of relative motion

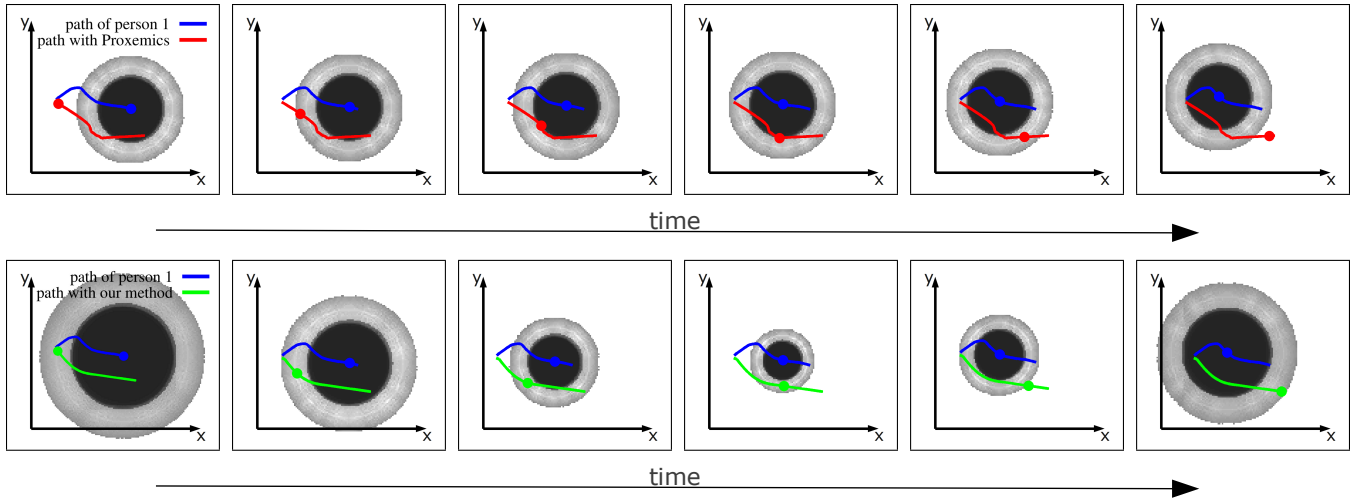


Fig. 4. Cost map evolution over time as generated by the fixed-distance baseline method (top) and our method (bottom). The darker the color the higher the costs. **Top:** The specified distance around the person does not change over time leading to an overly long avoidance maneuver around it. **Bottom:** The proposed technique based on relative motion prototypes produces a dynamic cost map whose high-cost radius changes over time producing a path that does not overshoot and better resembles the original human motion behavior in this situation.

prototypes composed of n elements that refer to n social contexts. An example of \mathcal{P} with five social contexts is shown in Fig. 5.

III. PLANNING WITH RMPs

Given n relative motion prototypes, this section presents the methods employed to derive socially-aware paths from the learned models. This includes the steps of social context detection, cost map computation and planning.

A. Social Context Detection

The recognition of the social context that enables the robot to select the appropriate prototype \mathcal{R}^* from \mathcal{P} is done in the same way than the model selection in the training phase described above. As soon as a Kalman filter-based person tracker delivers stable state estimates after initialization, we predict the trajectory into the future to determine the angle of approach α (see Fig. 3) and its corresponding model \mathcal{R}^* . The selected RMP is finally time-warped to match the velocity of the approaching subject.

B. Cost Map Computation and Path Planning

The selected motion prototype \mathcal{R}^* is then used to derive a dynamic cost map in a regular grid tessellation of the environment. This is done by computing a time-varying cost function for each grid cell that follows the learned distances $d_{i,j}(t)$ as specified by \mathcal{R}^* . As cost function we use a Gaussian distribution with mean $\mathbf{z}(t)$ and standard deviation $\sigma = d_{i,j}(t)/3$ at time t laid over the grid. This leads to smooth dynamic cost peaks around human subjects. The costs are updated in each step over the entire duration of the interaction (see Fig. 4).

Based on these costs, we employ a Θ^* planner [2] with a L_2 -norm heuristics. Θ^* is an any-angle A^* variant that is not constraint to paths formed by edges of the grid. This results in more natural and shorter any-angle paths with fewer heading changes than A^* -paths. We replan the path each time the cost map is updated.

IV. EXPERIMENTS

A. Data Set

We use the Edinburgh Informatics Forum Pedestrian Database (EIPD) [10] to learn the behavior of walking people. The data set has been collected with a stationary overhead camera at 23m height that observes a university hall of dimensions $15.8m \times 11.8m$ over a duration of several months. The data files contain annotated tracks of people defined as sequences of x and y coordinates of their center of gravity. A view of the scene is shown in Fig. 2.

B. Training

We have taken 90 pairwise interactions of people consisting in 180 tracks to learn the relative motion prototypes \mathcal{P} as described in Section II. As the provided annotations contain many double entries, missing or merged labels we had to manually inspect the trajectories to ensure that we learn only from correctly annotated examples. The BIC procedure returned a best value of $n = 5$ with angles of approach being 45° , 90° , 135° , and 180° . The DTW distance function renders matching invariant to track velocity but the case for zero velocity, that is when one person is standing still, required special treatment since the motion direction estimation is undefined. This yields the fifth prototype for which no angle of approach is defined.

	DW- A^*	DW- Θ^*	AD- A^*	AD- Θ^*
<i>Proxemics</i>	20.31	17.70	0.43m	0.39m
<i>RMP</i>	9.08	8.95	0.20m	0.21m
Improvement	55.3%	49.4%	53.5%	46.2%

TABLE I

COMPARISON OF SIMILARITY BETWEEN PLANNED (RMP) AND ORIGINAL PATH, USED AS THE SUBJECTIVE MEASURE. THE PROXEMICS-BASED BASELINE IS CLEARLY LESS ABLE TO REPRODUCE HUMAN RELATIVE MOTION.

	A^*	Θ^*
<i>Proxemics</i>	+36.1%	+29.0%
<i>RMP</i>	+14.7%	+11.9%
Improvement	59.3%	58.9%

TABLE II

COMPARISON OF RELATIVE PATH LENGTH USED AS THE OBJECTIVE MEASURE. THE PERCENTAGES ARE RELATIVE TO THE ORIGINAL PATH BY THE HUMAN SUBJECTS WHICH WOULD BE 0%. THE PROPOSED TECHNIQUE (RMP) IS CLEARLY MORE EFFICIENT THAN THE BASELINE WITH EVERY COMBINATION OF PLANNERS.

The learned prototypes in \mathcal{P} are shown in Fig. 5 together with the distances that mark the inner and outer border of the Personal space according to the Proxemics theory. It is interesting to note that intrusion into the Personal space by walking subjects appears to be common whereas for interaction such a nearness is accepted only for individuals of small social distance such as good friends or family members. The intrusion appears to be acceptable even if one of the subjects is standing. This is a finding that confirms our skepticism toward model-based approaches for the reasons listed in Section I and shows that a Proxemics-based avoidance strategy would lead to an overly conservative and inefficient robot motion behavior.

C. Quantitative Results

For testing we evaluate the learned models on a different set of 182 pairwise interactions. For each interaction we replace the path of one of the two persons (picked randomly) by a path generated with our system keeping the start and end positions fixed. We then quantify the performance of our system by computing two measures, an objective one – the path length – and a subjective one. As it is common to many HRI-related problems, we assume that social acceptability or comfort is highest when the robot exhibits the same behavior as a human. Accordingly, the subjective criterion is the aDTW distance (DW) which quantifies the similarity between the planned path and the original path of the human. We compare two variants of the proposed approach, one that makes use of Θ^* planning and another one that uses standard A^* . They both operate on a grid with square cells of size 0.05 meters. As a baseline, we replace the distance functions \mathcal{R} in the learned RMPs by a fixed

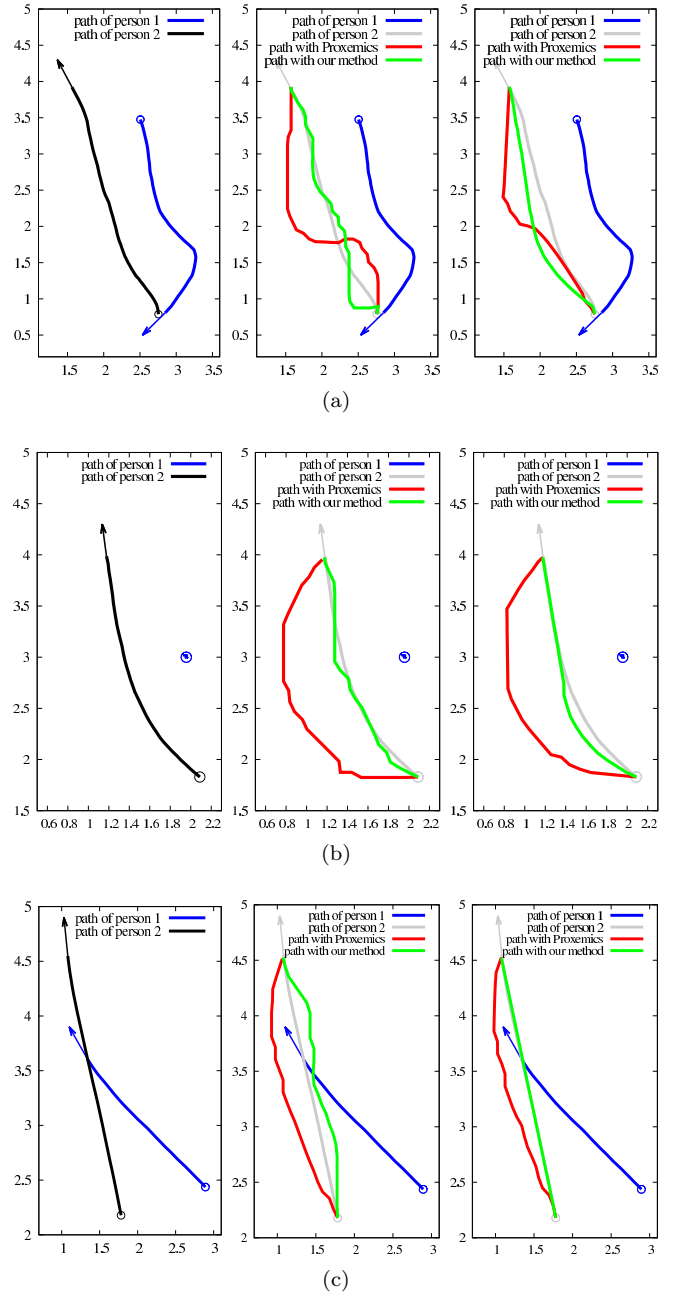


Fig. 6. Visualization of the planned paths using the learned relative motion prototypes and the paths originally taken by human subjects. The left column shows the original paths of the two persons as black and blue lines. The walking directions are marked by arrows. In the center and right columns, we replace the path of “person 2” by the Proxemics-based baseline method (red) and the planned path from our RMP approach (green). For planning either A^* (center) or Θ^* (right) is used. The images clearly show the advantage of our approach with respect to both objective (path length) and subjective criteria (similarity to human path).

distance d_{prox} derived from the Proxemics theory. This results in an avoidance behavior that seeks to keep the robot always d_{prox} meters away from humans. As d_{prox} we have taken the outer bound of the Personal space, that is, $d_{\text{prox}} = 2.0$ meters as the average distance of the Social space (see [4]).

The results are given in Table I. Low values indicate high similarity between the planned and the original paths. The proposed RMP approach clearly outperforms the Proxemics-based baseline by almost 50% with every combination of planners. To confirm that this result is not due to our modification of the DTW algorithm, we also compute another error metric, namely the average distance (AD) with respect to the closest point. The improvements are almost identical.

As the objective measure to evaluate the learned prototypes, we compute the relative length difference between the planned and the original path in percent. The results of the comparisons are shown in Table II. The baseline using a fixed distance performs modestly (up to 36% longer than the original trajectory) and results in longer paths than the learned RMPs whose length is only increased by 15%.

A visualization of three example interactions with varying angles of approach is shown in Fig. 6. It can be seen that the learned prototypes are able to reproduce the relative motion of humans to most extent while the baseline is overly conservative and contains more unnecessary heading changes.

V. CONCLUSIONS

In this paper we have addressed the problem of learning and reproducing motion behavior of walking people for the task of socially-aware robot motion planning among humans. Opposed to related work, we have taken a learning approach to the problem using observations of humans in publicly available large-scale surveillance data sets. The approach uses unsupervised learning to produce a set of Relative Motion Prototypes (RMP). RMPs are experimentally shown to optimize both an objective and a subjective performance measure and to decrease the errors by a factor of two. Further, paths generated with RMPs are similar to human paths and clearly outperform a Proxemics-based baseline method.

In future work, we plan to extend relative motion prototypes for multiple people and implement them on a real robot. RMPs can be elegantly combined by superposition of their individual cost maps and preliminary experiments with multiple persons have shown that the resulting motion behavior is natural and efficient. Furthermore, the current method runs in real-time but it is not computationally efficient since replanning occurs frequently with a dynamically updated cost map. We will thus consider incremental Theta* planners able to build on previous search results.

ACKNOWLEDGMENT

This work has been supported by the German Research Foundation (DFG) under contract number SFB/TR-8.

REFERENCES

- [1] S. Bitgood and S. Dukes, "Not another step! Economy of movement and pedestrian choice point behavior in shopping malls," *Environment and Behavior*, no. 3, pp. 394–405, 2006.
- [2] K. Daniel, A. Nash, S. Koenig, and A. Felner, "Theta*: Any-angle path planning on grids," *Journal of Artificial Intelligence Research (JAIR)*, vol. 39, pp. 533–579, 2010.
- [3] W. H. Day and H. Edelsbrunner, "Efficient algorithms for agglomerative hierarchical clustering methods," *Journal of Classification*, vol. 1, no. 1, pp. 7–24, 1984.
- [4] E. Hall, *Handbook of Proxemics Research*. Society for the Anthropology of Visual Communications, 1974.
- [5] J. M. Hammersley and P. Clifford, "Markov field on finite graphs and lattices," 1971.
- [6] D. Helbing and P. Molnar, "A social force model for pedestrian dynamics," *Phys. Rev. E*, vol. 51, pp. 4284–4286, 1995.
- [7] P. Henry, C. Vollmer, B. Ferris, and D. Fox, "Learning to navigate through crowded environments," in *Int. Conf. on Robotics and Automation (ICRA)*, Anchorage, USA, 2010.
- [8] J. D. Lafferty, A. McCallum, and F. C. N. Pereira, "Conditional random fields: Probabilistic models for segmenting and labeling sequence data," in *Int. Conf. on Machine Learning (ICML)*, 2001.
- [9] J. Mainprice, E. Sisbot, T. Simeon, and R. Alami, "Planning safe and legible hand-over motions for human-robot interaction," in *IARP Workshop on Technical Challenges for Dependable Robots in Human Environments*, Shanghai, China, 2010.
- [10] B. Majecka, "Statistical models of pedestrian behaviour in the forum," Master's thesis, School of Informatics, University of Edinburgh, 2009.
- [11] R. Middlemist, E. Knowles, and C. Matter, "Personal space invasions in the lavatory: suggestive evidence for arousal," *Journal of Personality and Social Psychology*, vol. 33, no. 5, pp. 541–6, 1976.
- [12] Y. Nakauchi and R. Simmons, "A social robot that stands in line," *Autonomous Robots*, vol. 12, no. 3, pp. 313–324, May 2002.
- [13] E. Pacchierotti, H. I. Christensen, and P. Jensfelt, "Embodied social interaction for service robots in hallway environments," in *Proc. of the Int. Conf. on Field and Service Robotics (FSR)*, 2005.
- [14] J. Rios-Martinez, C. Laugier, and A. Spalanzani, "Understanding human interaction for probabilistic autonomous navigation using risk-RRT approach," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, San Francisco, USA, 2011.
- [15] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE Trans. on Acoustics, Speech and Signal Processing*, vol. 26, no. 1, pp. 43 – 49, 1978.
- [16] L. Scandolo and T. Fraichard, "An anthropomorphic navigation scheme for dynamic scenarios," in *Int. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, February 2011.
- [17] Y. Tamura, T. Fukuzawa, and H. Asama, "Smooth collision avoidance in human-robot coexisting environments," in *Int. Conf. on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, 2010.