

Online Trajectory Prediction and Planning for Social Robot Navigation

Christoph Rösmann, Malte Oeljeklaus, Frank Hoffmann and Torsten Bertram

IEEE International Conference on Advanced Intelligent Mechatronics, Munich, Germany, 2017

“©2017 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Online Trajectory Prediction and Planning for Social Robot Navigation

Christoph Rösmann¹, Malte Oeljeklaus¹, Frank Hoffmann¹ and Torsten Bertram¹

Abstract—This paper addresses the safe and legible navigation of mobile robots in multi-agent encounters. A novel motion model provides the basis to predict, plan and coordinate agent trajectories in intersection scenarios. The approach establishes an implicit, non-overt cooperation between the robot and humans by linking the prediction and planning of agent trajectories within a unified representation in terms of timed elastic bands. The planning process maintains multiple topological alternatives to resolve the encounter in a manner compliant with the implicit rules and objectives of human proxemics. The trajectory is obtained by optimizing the timed elastic band considering multiple conflicting objectives such as fastest path and minimal spatial separation among agents but also global proxemic aspects such as motion coherence within a group. Cooperation is achieved by coupling predicted and planned agent trajectories to eventually reach an implicit agreement of the agents on how to circumnavigate each other. The parameters of the cost functions of the underlying motion model are identified by inverse optimal control from a dataset of 73 recorded encounters with up to five humans and a total of 283 individual trajectories. Playback simulations of recorded encounters and experiments with a robot traversing a group of oncoming humans demonstrate the feasibility of the approach to resolve general proxemic encounters.

I. INTRODUCTION

Mobile service robots provide communicative and informal services at public places, e.g. in museums, malls or fairs. In this context a legible and coherent proxemic interaction between humans and robots is a key factor for their consumer acceptance. Service robots are supposed to navigate in crowded environments in a way that is compliant with human proxemics. Such a proxemic perspective requires concepts that extend beyond the established methods for motion planning and mere obstacle avoidance.

This paper investigates proxemic aspects in the spatial interaction between humans and mobile robots. Proxemics investigates the mechanisms by which a human determines his pose and posture towards an interaction partner or its environment. Our research objective is to incorporate these proxemic aspects into the trajectory planning of mobile robots in indoor environments. Even though proxemics in general rests upon a multitude of cues such as pose, posture, gaze and gestures, we hypothesize that the overt motion is a sufficient cue to negotiate the traversal of agents in confined spaces. The conception and implementation of legible mobile robot proxemic behaviors involves the following steps: (i) recording of human proxemic encounter trajectory data with a visual tracker, (ii) online trajectory prediction and

planning with timed elastic bands (TEB), (iii) incorporation of proxemic objectives in the cost function and (iv) identification of the cost function parameters in the underlying motion model from recorded human-human encounters by means of inverse optimal control.

To find a safe and collision-free way through dynamic and dense crowds in real-time is a challenge known as the *freezing-robot-problem* (FRP) [1] in the literature. The uncertainty inherent to opponents future trajectories prevents a solution by classical global, offline motion planners [2]. In order to cope with dynamic scenarios elastic bands modify an initial global plan locally [3]. The extension to TEBs deforms trajectories rather than paths and thus unifies online motion planning and control [4]. The TEB planner is available online for the Robot Operating System (ROS).

Trautman [5] provides an approach that solves the FRP in a probabilistic framework. The author shows that an explicit agent cooperation is key to solving the FRP in case of group encounters. Gaussian-processes are adapted towards demonstrations in order to model human motion patterns. Sampled trajectories are evaluated by an interaction potential to incorporate spatial-temporal interactions among agents.

Human motion models for social robots are developed [6]–[8]. Proxemic aspects in lateral passing scenarios between humans and robots are considered in [9], [10]. [11] addresses similar scenarios using dynamic motion prototypes which are learned from human demonstrations.

A simulation-based motion model suited for pedestrian encounters is presented in [12]. The multi-hypothesis approach rests upon a deterministic energy minimization in terms of the predicted minimal distance between two agents. It considers multiple trajectories of alternative topologies. A sampling based approach for multi-robot collision avoidance is provided in [13]. Van den Berg et al. [14] extend velocity obstacles to support reciprocal collision avoidance. Kretzschmar et al. [15] introduce feature-based probabilistic models for multi-agent encounters. The human intention and objectives are captured by features such as time, acceleration, velocity and collision-avoidance. Spline-based trajectories are optimized with respect to alternative topological variants. Maximum entropy learning is applied to adapt the model to human demonstrations. The main difference to our contribution is the explicit coordination of topological variants, whereas our approach rests upon an implicit coordination in which human intentions are indirectly inferred from the perceived motion. Another feature-based method by [16] models the pedestrians spatial behavior as a Markov Decision Process. Inverse reinforcement learning is applied to estimate underlying feature weights.

¹Authors are with the faculty of Electrical Engineering & Information Technology, Institute of Control Theory and Systems Engineering, TU Dortmund University, 44227 Dortmund, Germany {forename.surname}@tu-dortmund.de

II. LOCAL TRAJECTORY OPTIMIZATION

A. Constrained Optimization Problem

The TEB approach presented in [4], [17] optimizes trajectories by subsequent modification of an initial trajectory generated by a global planner. A trajectory $\mathcal{B} = \bigcup_{k=1}^n \{\mathbf{s}_k, \Delta T_k\} \setminus \{\Delta T_n\}$ is represented by an ordered sequence of n poses augmented with $n - 1$ time intervals. $\mathbf{s}_k = [x_k, y_k, \beta_k]^\top \in \mathbb{R}^2 \times S^1$ denotes the pose of the robot at time stamp k and $\Delta T_k \in \mathbb{R}^+$ represents the time interval associated with the transition between consecutive poses \mathbf{s}_k and \mathbf{s}_{k+1} . The time-optimal trajectory \mathcal{B}^* is obtained by solving the nonlinear program (NLP):

$$\min_{\mathcal{B}} \sum_{k=1}^{n-1} \Delta T_k^2 \quad (\text{NLP})$$

subject to:

$$\begin{aligned} \mathbf{s}_1 &= \mathbf{s}_s, \quad \mathbf{s}_n = \mathbf{s}_f, \quad 0 \leq \Delta T_k \leq \Delta T_{\max}, \\ \mathbf{h}_k(\mathbf{s}_{k+1}, \mathbf{s}_k) &= \mathbf{0}, & (\text{Non-holonomic kinematics}) \\ \mathbf{o}_k(\mathbf{s}_k) &\geq \mathbf{0}, & (\text{Clearance from obstacles}) \\ \boldsymbol{\nu}_k(\mathbf{s}_{k+1}, \mathbf{s}_k, \Delta T_k) &\geq \mathbf{0}, & (\text{Velocity limits}) \\ \boldsymbol{\alpha}_k(\mathbf{s}_{k+1}, \mathbf{s}_k, \mathbf{s}_{k-1}, \Delta T_{k+1}, \Delta T_k) &\geq \mathbf{0}. & (\text{Acceleration limits}) \end{aligned}$$

Initial and final poses, \mathbf{s}_1 and \mathbf{s}_n , are tied with the current robot pose \mathbf{s}_s obtained from robot localization and goal pose \mathbf{s}_f respectively. ΔT_{\max} provides an upper bound on ΔT_k in order to accomplish an appropriate discretization of the continuous time motion. Equality and inequality equations restrict the feasible set w.r.t. environmental and robotic constraints such as non-holonomic kinematics, clearance from obstacles and bounds on velocities and accelerations. A detailed discussion of the constraints is beyond the scope of this paper, the interested reader is referred to [17]. Note, customized constraints might be added to (NLP) for particular robots and applications.

B. Approximative Least-Squares Optimization

Instead of solving (NLP) directly, it is mapped into an unconstrained least-squares optimization problem with constraints approximation. Consequently, efficient nonlinear least-squares solvers which approximate the Hessian by first order derivatives and exploit the sparsity pattern of the problem are utilized [17]. In the following the arguments of constraints are omitted for the sake of readability. The approximated optimization problem is defined by:

$$\begin{aligned} \mathcal{B}^* &= \arg \min_{\mathcal{B} \setminus \{\mathbf{s}_1, \mathbf{s}_n\}} V(\mathcal{B}) \quad (1) \\ V(\mathcal{B}) &= \sum_{k=1}^{n-1} \left[\Delta T_k^2 + \sigma_h \|\mathbf{h}_k\|_2^2 + \sigma_\nu \|\min\{\mathbf{0}, \boldsymbol{\nu}_k\}\|_2^2 + \dots \right. \\ &\quad \left. + \sigma_o \|\min\{\mathbf{0}, \mathbf{o}_k\}\|_2^2 + \sigma_\alpha \|\min\{\mathbf{0}, \boldsymbol{\alpha}_k\}\|_2^2 \right] \\ &= \mathbf{w}^\top \mathbf{f}(\mathcal{B}) \end{aligned}$$

Constraints of (NLP) are expressed in terms of quadratic penalties with weights σ . For inequalities respectively one-sided penalty terms the min-operator applies row-wise. Constraints for \mathbf{s}_s and \mathbf{s}_f are eliminated by substitution and are

therefore not subject to the optimization. For the remainder, cost function $V(\mathcal{B})$ is expressed in terms of the dot product, in which \mathbf{w} captures individual weights and $\mathbf{f}(\mathcal{B})$ contains individual cost terms. The optimization problem (1) is efficiently solved with a sparse variant of the Levenberg-Marquardt (LM) Algorithm [17].

Note, optimization is integrated with state feedback to repeatedly refine the previous solution and in order to react on new way-points of the global plan, dynamic environments and disturbances. The resulting predictive control strategy mimics a receding horizon until the global goal is reached. In each online invocation of the optimization procedure the number of poses n is increased or decreased by comparing current time intervals ΔT_k with a desired ΔT_{ref} . Hence allowing the decoupling of remaining path length and overall transition time while retaining the imposed temporal discretization [17].

III. IMPLICITLY COORDINATED MOTION MODEL

In [17] the TEB is extended to the exploration and optimization of multiple candidate trajectories with quasi-static obstacles. This approach is further extended to predict and incorporate dynamic motions of other agents. The prediction of other agents future trajectories is essential to achieve a truly proxemic behavior. The implicitly coordinated motion model (ICMM) captures proxemic aspects of spatial-temporal interaction and thus provides the basis to predict the agents trajectories as well as to plan the robots own motion. The most likely resolution of a proxemic encounter is identified in three stages. First, a set of topological alternatives are identified from that plausible solutions are selected. Second, locally optimal trajectories for all alternative topologies are planned in parallel by the TEB optimization. Finally, the least-cost trajectory is selected. Parallel planning is performed for all agents. In the case of humans the selected trajectory provides a prediction. In case of the robot the first translational and rotational velocity of the selected trajectory are commanded according to the predictive control strategy. Assuming that the agents act in a cooperative manner, the trajectories mutually effect each other until the agents implicitly achieve a common resolution.

A. Exploration and Optimization in Distinctive Topologies

Encounters with multiple agents allow alternative topological trajectories in which the individuals pass each other either to the left or right. The intentions of agents reveal themselves by their overt motion upon approach. For solving (1) efficiently during runtime, local optimization schemes (such as LM) are utilized. Therefore, the local solution depends on the initialization and the optimizer is unable to transit to a different local optimal trajectory with agents passing on the opposite side.

In order to distinguish whether two paths resp. trajectory belong to the same topology, an equivalence relation (invariant) based on the concept of complex analysis is borrowed from [18]. The invariant $\mathcal{H}(\tau) = \int_\tau \mathcal{F}(z) dz$ with $z = x + iy \in \mathbb{C}$ of a path τ is denoted as H -signature

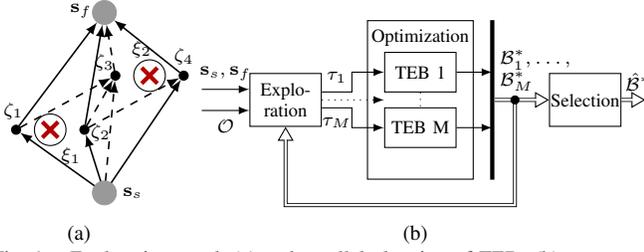


Fig. 1. Exploration graph (a) and parallel planning of TEBs (b)

and is equal for trajectories of the same topology. The integrand \mathcal{F} depends on the location of obstacles in the plane. In the following obstacles are modeled by connected domains $\mathcal{O}_1, \mathcal{O}_2, \dots, \mathcal{O}_N$. For each obstacle, $\xi_l \in \mathcal{O}_l, \forall l = 1, 2, \dots, N$, denotes a *representative point* in a complex representation. Given an arbitrary analytic function f_0 over \mathbb{C} , the *obstacle marker function* \mathcal{F} is defined by:

$$\mathcal{F}(z) = \frac{f_0(z)}{(z - \xi_1)(z - \xi_2) \cdots (z - \xi_N)} \quad (2)$$

Established f_0 terms and the analytic solution of the H -signature for discrete trajectories like \mathcal{B} are provided in [17].

Agents might pass each other either to the left or right. For each agent an exploration graph is generated by considering all permutations of pairwise left and right passings (see Fig. 1(a)). Let s_s denote the current pose of an agent and s_f its designated goal pose. ξ_l represents the pose of the l -th interacting agent. Initial collision free paths of alternative topologies are generated only considering the agents that constitute possible obstacles for a path that connects s_s and s_f : (i) For each agent add a pair of nodes ζ_i to the left and right side to the graph. (ii) Connect nodes by forward-directed edges which orientation is sufficiently similar (here: $\Delta\theta \leq \pi/4$) to the straight line from s_s and s_f . (iii) Extract all primitive paths from the resulting acyclic graph using depths-first search (DFS). (iv) Calculate the H -signature for each path and remove redundant paths that possess the same H -signature.

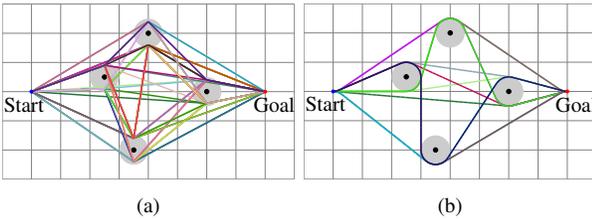


Fig. 2. Exploration example with four agents between start and goal, (a) feasible paths, (b) paths after homotopy filtering and TEB optimization.

The exploration stage generates M initial candidate paths that belong to alternative topologies. Those are converted into TEBs by distributing initial poses uniformly along the path. The TEBs are optimized in parallel (see Fig. 1(b)) with respect to the cost function in (1). In order to select a proper TEB for an agent in the current sampling interval, the original cost of (1) is augmented by criteria that capture the global proxemic aspects of the entire encounter (only for selection, not optimization) namely

- Spatial-temporal integrity of agents that move in a group: A penalty is imposed for trajectories that pen-

etrate a group. Permanent groups are detected by a modified DBSCAN [19] which clusters agents in time and space according to their coherent motion.

- Right versus left bias: In case of right left symmetric trajectories of otherwise equal costs this offset induces a bias, for example to prefer right hand traffic.
- Initial curvature: Trajectories which deviate less from the straight line to the goal are preferred. The curvature of the TEB along the first H nodes is computed from the absolute change of orientation $1/H \sum_{i=1}^H |\beta_{i+1} - \beta_i|$.

The traditional and proxemic objectives are expressed by an augmented cost function $V_c(\mathcal{B}) = \mathbf{w}_c^T \mathbf{f}_c(\mathcal{B})$ and the optimal compromise solution is determined by their relative weights \mathbf{w}_c . The best TEB $\hat{\mathcal{B}}^*$ is obtained by solving:

$$\hat{\mathcal{B}}^* = \arg \min_{\mathcal{B}_p^* \in \{\mathcal{B}_1^*, \mathcal{B}_2^*, \dots, \mathcal{B}_M^*\}} V_c(\mathcal{B}_p^*) \quad (3)$$

It is impossible to quantify the weights \mathbf{w} and \mathbf{w}_c in advance as the relationship between a desired proxemic behavior and the weights is nontransparent. Instead the weights are identified by learning from demonstration in section IV, such that the optimal solutions (3) coincide with the trajectories recorded in human encounters.

Fig. 2 shows an example of the exploration with four additional agents. The initial DFS generates 52 paths that are compliant with the forward condition. From this set only twelve topological different paths remain after filtering redundant solutions according to the H -signature (see Fig. 2(b)). The computation of the graph requires less than 1 ms on a common desktop computer and the optimization of each trajectory approx. 5–10 ms (single-threaded). Note, this exploration strategy assumes relatively small agent footprints which is usually fulfilled for humans. For larger shapes a sampling based strategy from [17] might be adapted.

The ongoing motion of agents requires an update of the TEBs and their H -signatures, the elimination of obsolete TEBs or an instantiation of novel TEBs. Thus, the exploration stage (see Fig. 1(b)) is extended to: (i) verify TEBs for compliance with the forward condition (ii) discard non-compliant obsolete TEBs (iii) the initial and final nodes of the TEBs are updated with the agents current start and goal pose (iv) redundant TEBs with identical H -signatures are removed (v) novel TEBs are instantiated for homotopy classes that are not covered by the current set of TEBs (vi) optimize the TEBs according to (1).

B. Cooperative Planning

The above TEB framework enables an agent to plan his trajectory under the assumption that the trajectories of the interacting agents are known. There is no explicit communication among agents, but they indirectly signal their intentions as they yield to either side. This signaling triggers a mirrored movement of the partner agent, which eventually leads to an implicit mutual agreement. The whole encounter is resolved by coupling the agents predicted trajectories. Notice the correlation between an agents decision and his prediction \hat{B}_i^* of the movement of other agents. Fig. 3 shows

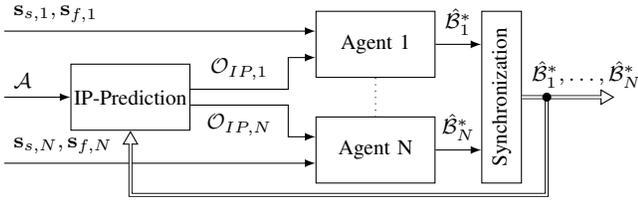


Fig. 3. Planning parallel TEBs in dynamic environments

the overall scheme of the implicitly coordinated motion planner. The set of agents is denoted by \mathcal{A} . A subscript index for each agent is added to start pose, final pose and trajectories respectively. Each agent maintains its own subset of feasible trajectories in alternative topologies from which it selects the TEB of least proxemic cost (see Fig. 1(b)). From the predicted trajectories of each agent pair an *interaction point* (IP) is computed from the two corresponding nodes of minimal spatial-temporal separation among both trajectories. Each agent temporarily considers his IPs with other agents as static obstacles \mathcal{O}_{IP} in its current local TEB optimization. Upon recomputation of IPs an agent refines its current preferred TEB. An alternative topology might become more desirable in case of a substantial change in IP locations. Switching to a topologically alternative trajectory might cause other agents to shift their path. It turns out, that in practice such a cascade of shifts in the agents topology ebbs away within one to three iterations. Upon the synchronization of agent trajectories in Fig. 3 the start and goal poses of the TEBs are updated according to observed motion of the agents. As the TEB approach is an online planner novel observations are immediately incorporated to refine the TEB.

The ICMM enables robot motion planning and control by considering the robot as just another agent with the same underlying proxemic motion model estimated from human demonstrations. Rather than to predict the robots TEB plans the robots trajectory. The robot motion model might differ in terms of dynamic and kinematic constraints of the platform. In our case the maximum robots translational velocity of about 0.8 m/s is lower than the typical human walking speed of ≈ 1.5 m/s. The ICMM assumes a goal pose for each agent, which in our case is predicted by a modified version of virtual goals presented in [20]. With a constant horizon of planning the goal is assumed 5 m ahead of the agent with the heading determined by the initial direction of agent motion averaged over the first few steps. The goal heading is redefined in case the agents current heading deviates by more than $\delta\theta = \pi/3$ from the original goal heading.

IV. LEARNING FROM DEMONSTRATION

As it is nearly impossible to quantify the weights \mathbf{w}_c in the cost function in (3) its underlying parameters are learned from recorded trajectories in encounters of humans in groups. For that purpose people in an indoor room are tracked by a birds eye view camera. The subjects wear colored t-shirts to facilitate their segmentation and registration in the birds eye view. The camera is calibrated such that the location of people in the image is mapped to their planar 2D pose in the room. The raw trajectories are filtered and smoothed

by fitting them to B-splines. The recorded dataset includes 73 demonstrations of encounters with up to five interacting agents with a total of 283 individual trajectories. The particular encounters included head on as well as diagonal trajectories of two groups with 2 vs. 1, 3 vs. 1, 2 vs. 2, 3 vs. 2 agents. The duration of scenarios varied between 5 – 10 s during which the subjects covered distances between 6 – 9 m.

The parameter identification includes the parameters of the local costs functions, e.g. comfortable spatial separation, as well as the global aspects such as group integrity. The TEB cost function in (1) contains the individual objective functions and constraints. (i) The minimal spatial separation of two agents at the closest encounter is obtained by averaging over the separation observed in the 73 scenarios. (ii) The maximum translational velocity of human agents is set to the average velocity across all demonstrations. (iii) The maximum rotational velocities and acceleration bounds are set to the average of maximum values in each trajectory.

The weights \mathbf{w}^* that determine the relative importance of conflicting objectives are identified by a subgradient optimization technique [21], such that the optimal trajectories according to the model coincide with the actual recorded trajectories. The optimal weights are those which minimize the discrepancy in cost function between the demonstrated and planned trajectory. Assuming, that the optimal TEB trajectory constitutes the global optimum of (1), the learning problem is embedded into a superior optimization problem:

$$\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathbb{R}^+} \underbrace{\frac{1}{D} \sum_{d=1}^D [\mathbf{w}^T \mathbf{f}(\check{\mathcal{B}}_d) - \min_{\mathcal{B}} \mathbf{w}^T \mathbf{f}(\mathcal{B})]}_{C(\mathbf{w})} \quad (4)$$

in which $\check{\mathcal{B}}_d$ denotes the $d = 1, 2, \dots, D$ demonstrated trajectories. The last term of (4) constitutes the optimum of the forward optimization (1). According to the maximum-margin-principle, (4) is augmented by a margin and a regularization term (refer to [21]). Due to the min-operator (4) is non-differentiable such that standard optimization techniques are not applicable. However, [22] proves that the subgradient method is able to solve convex optimization problems. The subgradient of $C(\mathbf{w})$ in (4) is obtained by

$$\partial_{\mathbf{w}} C(\mathbf{w}) = \frac{1}{D} \sum_{d=1}^D [\mathbf{f}(\check{\mathcal{B}}_d) - \mathbf{f}(\mathcal{B}^*)] \quad (5)$$

The superior optimization problem (4) is solved by subgradient descent:

$$\mathbf{w}_{j+1} = \mathbf{w}_j - \alpha_b \cdot \partial_{\mathbf{w}_j} C(\mathbf{w}_j) \quad (6)$$

α_b denotes the constant step size for each update.

The learning stage identifies weights of the proxemic terms (group penetration, left/right offset and curvature) in (3) for selecting the *best* trajectory. The topologies of trajectories $\mathcal{B}_1^*, \mathcal{B}_2^*, \dots, \mathcal{B}_M^*$ are matched with the demonstration by comparing their H -signatures. The trajectory that matches the demonstrations topology is denoted by $\check{\mathcal{B}}^*$. For learning, three types of events are defined at which H -signatures are

matched with the demonstration and the cost function (3) is evaluated: i) initial solution after the first iteration, ii) switching of the preferred topology in (3), iii) critical time-to-collision (the analysis of the recorded encounters indicates that humans signal their intent by altering their overt heading at approximated 1.4 s prior to the closest approach).

For $d = 1, 2, \dots, D$ demonstrations with E_d events each, let \mathcal{I} denote the set of indices of all recorded trajectories. The superior parameter optimization problem for (3) is given by:

$$\begin{aligned} C_c(\mathbf{w}_c) &= \frac{1}{N_{DE}} \sum_{i \in \mathcal{I}} \left[\mathbf{w}_c^T \mathbf{f}_c(\check{\mathcal{B}}_i^*) - \min_{\mathcal{B}_{b,i}^*} \mathbf{w}_c^T \mathbf{f}_c(\mathcal{B}_{b,i}^*) \right] \\ \mathbf{w}_c^* &= \arg \min_{\mathbf{w}_c \in \mathbb{R}} C_c(\mathbf{w}_c) \end{aligned} \quad (7)$$

with $\mathcal{B}_{b,i}^* \in \{\mathcal{B}_{1,i}^*, \mathcal{B}_{2,i}^*, \dots, \mathcal{B}_{M,i}^*\}$ as candidates in alternative topologies. N_{DE} denotes the number of elements in \mathcal{I} .

The structure of the parameter identification problem (7) is similar to (4), with the difference that the min-operator selects among a discrete rather than quasi-continuous set of alternative trajectories. The solution is obtained by subgradient optimization with a subgradient of (7) defined by:

$$\partial_{\mathbf{w}_c} C_c(\mathbf{w}_c) = \frac{1}{N_{DE}} \sum_{i \in \mathcal{I}} \left[\mathbf{f}_c(\check{\mathcal{B}}_i^*) - \mathbf{f}_c(\hat{\mathcal{B}}_i^*) \right] \quad (8)$$

Hereby, $\hat{\mathcal{B}}_i^*$ minimizes $\mathbf{w}_c^T \mathbf{f}_c(\mathcal{B}_{b,i}^*)$. For a constant step size α_c , the subgradient descent update rule is given by:

$$\mathbf{w}_{c,j+1} = \mathbf{w}_{c,j} - \alpha_c \cdot \partial_{\mathbf{w}_{c,j}} C_c(\mathbf{w}_{c,j}) \quad (9)$$

V. EXPERIMENTAL RESULTS

This section investigates the learning of parameters of the proxemics motion model acquired from recorded demonstrations and analyzes simulated playbacks of actual encounters as well as closed loop experiments of a mobile robot (Pioneer 3DX) traversing groups of humans.

Fig. 4 shows the progress of the average trajectory error $\bar{\delta}_{Tr}(\mathbf{w}_{b,j})$, in terms of the mean distance between demonstrated and learned trajectory w.r.t. the iterations of subgradient descent. $\bar{\delta}_{Tr}(\mathbf{w}_{b,j})$ converges after approximately 15 iterations with a mean trajectory error of about 38.2 cm.

Fig. 5 shows the same evolution in terms of the classification rate of the correct homotopy class across all events. The non-monotonic improvement in event recall is typical for non-differentiable subgradient solutions. For the optimal weights the selection error decreases to 0.7%.

The analysis of the global proxemics parameters reveal a preference of agents to pass each other at the right hand side. The average separation at closest encounter between two interaction partners is 0.64 m. According to proxemics

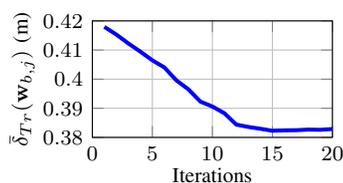


Fig. 4. Learning TEB cost function

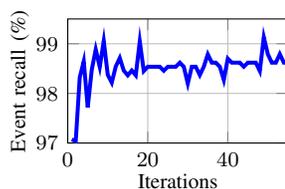


Fig. 5. Learning TEB selection

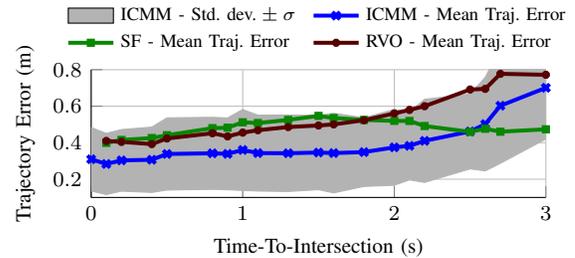


Fig. 6. Trajectory analysis on the entire dataset

[23], the social space extends to 1.20 m, whereas the distance of 0.64 m denotes a humans personal space. However, these proxemic distances correspond to static and permanent spatial configurations and the literature provides substantial evidence that for transient configurations humans are comfortable with substantially smaller separations [11].

In the following playback simulations and experiments on the basis of the ICMM are presented. In a playback simulation, the previously recorded human trajectories are fed back step by step into the ICMM, which then predicts the agents future trajectories. The ICMM prediction is compared with the future ground truth evolution of the encounter.

Fig. 6 shows the mean trajectory error and the standard deviation of the predicted trajectories as a function of the remaining time to closest approach. For comparison, the mean trajectory error of the *Social Forces* approach (SF) [20] and the *Reciprocal Velocity Obstacles* method (RVO) [24] determined on the very same dataset are included. Parameters for the minimum distance to other agents and the maximum velocity are chosen according to ICMM. ICMM predictions are slightly more precise for remaining times below 2.5 s.

Two snapshots of a 2 by 2 encounter are shown in Fig. 7. Initially, the ICMM predicts that agents 1 and 2 jointly pass agents 3 and 4 on their right side without penetrating their grouping. A brief instance later both groups split up their formation such that the resolution with shorter trajectories but dissociation of the groups is preferred. The ICMM switches the predicted homotopy class (see 7(b)). A conventional motion planner does not consider the cooperation of the other agents thus plans the safer detour to avoid the entire group.

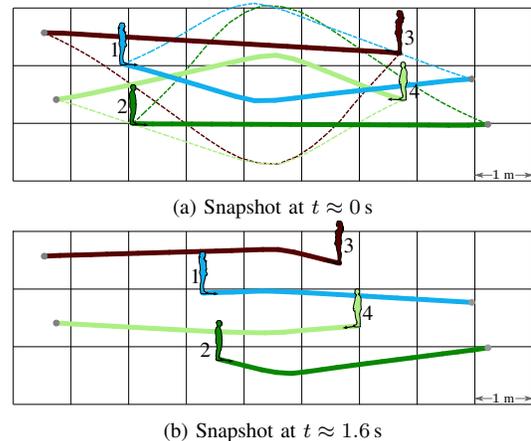


Fig. 7. 2 by 2 encounter: Solid lines represent predicted trajectories, dashed lines denote feasible alternative homotopy classes

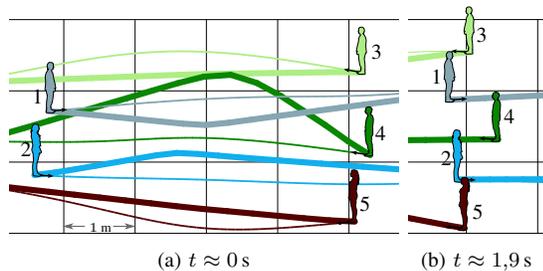


Fig. 8. 2 by 3 encounter: Solid lines represent predicted trajectories, dashed lines are recorded trajectories (demonstrations)

Fig. 8 illustrates a two by three encounter in which the homotopy class of predicted trajectories matches the actual resolution of the encounter fairly late. At time $t = 0$ s the predicted trajectories maintain the integrity of agents 1 and 2 assuming a cooperation of agent 4 by passing the group to his right. The dashed lines denote the actual trajectories in which agent 4 continues along his original heading thereby forcing agents 1 and 2 to split up. At time $t = 1.9$ s the predicted and actual trajectories coincide.

In a total of 43 experiments the robot traversed a group of humans without any collision. In the scenario in Fig. 9 the robot approaches three agents, again the actual trajectories are denoted by thin lines and predictions by thick lines. The predicted homotopy class of the encounter is revoked multiple times in response to the motion of agents 2 and 3. Initially at time $t = 2$ s the robot plans to traverse through the gap between agents 2 and 3 in Fig. 9(a). As agent 2 turns to the left the ICMM switches to the resolution in Fig. 9(b). Since the heading of agent 3 changes significantly its goal pose is reset to the north in Fig. 9(c) assuming that he passes the robot on his right side. In Fig. 9(c) the agent 3 turns once more to the right such that the robot reverts to the original plan which coincides with the actual outcome.

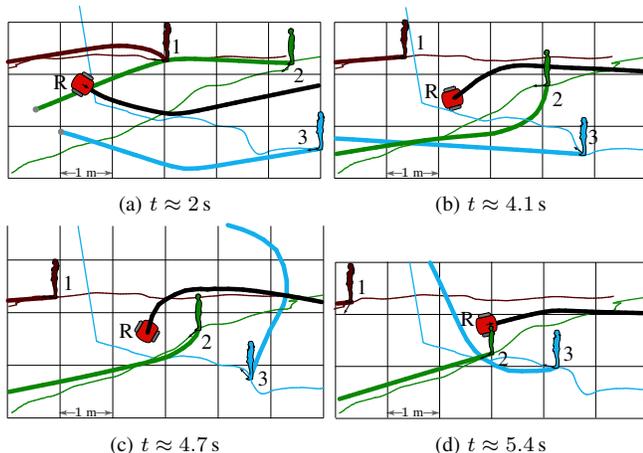


Fig. 9. Scenario with a real mobile robot

VI. CONCLUSIONS

This paper presents a novel approach for robot motion planning in crowds. The planning and prediction of trajectories rest upon a proxemic motion model which underlying cost functions are extracted from recorded human encounters.

The motion model captures proxemic aspects such as group integrity and right versus left biases. Implicit cooperation of agents is established by coupling planning and prediction in a unified representation with timed elastic bands. The major contribution is the ability to handle multiple topologically alternative resolutions of the encounter. This property allows the robot agent to adapt its plan in response to changes of the movements and intentions of opponents. Resimulation of recorded encounters and experiments with a robot demonstrate the feasibility of the approach to model, predict and plan trajectories in realistic proxemic encounters online.

REFERENCES

- [1] L. Blackmore, "Robust path planning and feedback design under stochastic uncertainty," in *AIAA Guidance, Navigation and Control Conference*, 2008.
- [2] S. M. LaValle, *Planning Algorithms*. New York, USA: Cambridge University Press, 2006.
- [3] S. Quinlan, *Real-Time Modification of Collision-Free Paths*. Stanford, USA: Stanford University, Department of Computer Science, 1995.
- [4] C. Rösmann, *et al.*, "Trajectory modification considering dynamic constraints of autonomous robots," in *7th German Conference on Robotics*, May 2012, pp. 74–79.
- [5] P. Trautman, "Robot navigation in dense crowds: statistical models and experimental studies of human robot cooperation," Ph.D. dissertation, California Institute of Technology, 2012.
- [6] M. Bennewitz, *et al.*, "Learning motion patterns of people for compliant robot motion," *Journal of Robotics Research*, vol. 24, no. 1, 2005.
- [7] G. Arechavaleta, *et al.*, "An optimality principle governing human walking," *IEEE Trans. on Robotics*, vol. 24, no. 1, pp. 5–14, 2008.
- [8] S. Thompson, *et al.*, "A probabilistic model of human motion and navigation intent for mobile robot path planning," in *IEEE 4th Int. Conf. on Autonomous Robots and Agents*, 2009.
- [9] E. Pacchierotti, *et al.*, "Embodied social interaction for service robots in hallway environments," in *5th Int. Conf. on Field and Service Robotics*. IEEE, 2005, pp. 476–487.
- [10] R. Kirby, "Social robot navigation," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 2010.
- [11] M. Luber, *et al.*, "Socially-aware robot navigation: A learning approach," in *Int. Conf. on Intelligent Robots and Systems*, 2012.
- [12] S. Pellegrini, *et al.*, "Wrong turn - no dead end: A stochastic pedestrian motion model," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 2010, pp. 15–22.
- [13] R. A. Knepper and D. Rus, "Pedestrian-inspired sampling-based multi-robot collision avoidance," in *RO-MAN*, 2012, pp. 94–100.
- [14] J. van den Berg, *et al.*, "Reciprocal n -body collision avoidance," in *14th Int. Symp. on Robotics Research*, 2009, pp. 3–19.
- [15] H. Kretzschmar, *et al.*, "Socially compliant mobile robot navigation via inverse reinforcement learning," *Int. J. Robot. Res.*, 2016.
- [16] S.-Y. Chung and H.-P. Huang, "Incremental learning of human social behaviors with feature-based spatial effects," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2012, pp. 2417–2422.
- [17] C. Rösmann, *et al.*, "Integrated online trajectory planning and optimization in distinctive topologies," *Robotics and Autonomous Systems*, vol. 88, pp. 142–153, 2017.
- [18] S. Bhattacharya, *et al.*, "Search-based path planning with homotopy class constraints," in *National Conf. on Artificial Intelligence*, 2010.
- [19] M. Ester, *et al.*, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *Int. Conf. on Knowledge Discovery and Data Mining*. AAAI Press, 1996, pp. 226–231.
- [20] M. Luber, *et al.*, "People tracking with human motion predictions from social forces," in *Int. Conf. on Robotics & Automation*, USA, 2010.
- [21] N. Ratliff, *et al.*, "Maximum margin planning," in *Int. Conf. on Machine Learning*, Juli 2006.
- [22] N. Shor, *Minimization methods for non-differentiable functions*, ser. Springer series in computational mathematics. Springer-Verlag, 1985.
- [23] E. T. Hall, "A system for the notation of proxemic behavior," *American Anthropologist*, vol. 65, no. 5, pp. 1003–1026, 1963.
- [24] J. van den Berg, *et al.*, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *IEEE Int. Conf. on Robotics and Automation*, May 2008, pp. 1928–1935.