

An Approach to Socially Compliant Leader Following for Mobile Robots

Markus Kuderer and Wolfram Burgard

Department of Computer Science, University of Freiburg

Abstract. Mobile robots are envisioned to provide more and more services in a shared environment with humans. A wide range of such tasks demand that the robot follows a human leader, including robotic co-workers in factories, autonomous shopping carts or robotic wheelchairs that autonomously navigate next to an accompanying pedestrian. Many authors proposed follow-the-leader approaches for mobile robots, which have also been applied to the problem of following pedestrians. Most of these approaches use local control methods to keep the robot at the desired position. However, they typically do not incorporate information about the natural navigation behavior of humans, who strongly interact with their environment. In this paper, we propose a learned, predictive model of interactive navigation behavior that enables a mobile robot to predict the trajectory of its leader and to compute a far-sighted plan that keeps the robot at its desired relative position. Extensive experiments in simulation as well as with a real robotic wheelchair suggest that our method outperforms state-of-the-art methods for following a human leader in wide variety of situations.

1 Introduction

There is a wide range of applications for mobile robots for which it is desirable that the robot follows a human leader. For example a robotic co-worker that provides tools to a human in a factory needs to stay in a position where the human can reach the robot. Similarly, a mobile shopping cart should always stay in a position where the human is able to place objects into it. A further application is a robotic wheelchair that stays side by side to an accompanying pedestrian, allowing interaction with the pedestrian during the navigation task.

When following a human leader, it is beneficial for the robot to reason about the natural navigation behavior of pedestrians. During navigation, pedestrians interact with their environment, which includes obstacles, other nearby humans and also the robot itself. A robot that has a better understanding of this interactive behavior is able to fulfill its task in a socially compliant way, i.e., in a way that does not unnecessarily hinder nearby pedestrians. Such a robot is able to predict the behavior of the humans and to plan far-sighted trajectories that keep the robot close to its desired position in the long run.

There has been a wide range of research on controlling a group of robots in formation, which have, to some extent, also been applied in the context of social

robotics [12, 14]. Many of these approaches utilize control-theoretic methods to steer the robot towards a virtual target that moves along with the leader [5, 14]. However, these methods mostly neglect information about the more complex navigation behavior of pedestrians that strongly depends on the environment.

In this paper, we propose to utilize a feature-based model of human navigation behavior to predict the path of the leading pedestrian [9]. This model accounts for the intention of a human to reach a certain goal while keeping a comfortable velocity, avoiding strong accelerations and to stay clear of obstacles. The individual characteristics of different pedestrians, or distinct behavior in different environments can be learned from observation.

The contribution of this paper is a method that simultaneously predicts the most likely trajectory of the pedestrian and computes the trajectory for the robot that minimizes the distance to its desired relative position along the whole trajectory in a forward-looking manner. Such a predictive planning method leads to a socially more compliant behavior of the robot. In addition, planning long-term trajectories mitigates the problem of local minima in a local control function, especially in the presence of arbitrary, non-convex obstacles in the environment. We conducted a simulated comparison of our method to related approaches as well as experiments with a real robot that show the applicability of the proposed approach to navigate a robotic wheelchair next to an accompanying pedestrian.

2 Related Work

In the past, many authors proposed methods to navigate a group of robots in formation. Liu et al. [11] cast the joint path planning task of a robot formation as a linear programming problem. Similar to our approach, they plan the trajectories to the target position of each robot. However, Liu et al. control the group of robots in a central manner and each robot executes the optimal trajectory. Balch and Hybinette [1] propose to use social potential fields that pull the robots towards attraction points to achieve a certain formation. Our experiments include a comparison to a social potential-based approach.

A different problem arises when the task of the robot is to follow a leader whose goal is unknown. Chiem and Cervera [4] and Huang et al. [8] propose to compute a cubic Beziér curve between the leading robot and the follower. The follower then navigates along this trajectory, using a velocity controller. In addition, if the robots task is not only to follow the same path but to stay in a certain formation, they propose to compute virtual targets for each of the robots and compute Beziér curves to these target positions. However, they follow the leading robot without active obstacle avoidance. Desai et al. [6] and Das et al. [5] use control theoretic approaches to keep each robot close to its designated position within the formation, also considering obstacles in the environment. If the desired shape of the formation changes, they introduce control graphs to assign the robots to their new position in the formation. Qin et al. [15] use artificial forces to navigate each robot close to the desired position in a formation. Similarly, Tanner and Kumar [16] propose to use navigation functions to keep a

group of robots in a certain formation. Navigation functions also lead the robot along the gradient of a smooth function, similar to artificial forces, but there are no local minima allowed, except of the target position. In general, however, it is difficult to design such a function for arbitrary environments [10]. Chen and Wang [3] provide a survey on different approaches to robot formation control.

The abovementioned approaches use local control methods to steer the robot either directly to the desired position in the formation, or to some local virtual target position. In contrast, we predict the trajectory of the leader based on its current state and the state of the environment. At the same time, we compute the trajectory that minimizes the distance to the desired relative position along this trajectory while satisfying further constraints. This prevents the robot to get stuck in local minima of the cost function and allows it to adapt the planned trajectories to the environment early on.

Similar methods have also been used to enable a robot to follow a human leader. Pradhan et al. [13] utilize a navigation function method and set the tracked positions of the pedestrian as virtual target positions. Therefore, the robot is only able to follow the person, but not to stay at a fixed relative position. Prassler et al. [14] aim at coordinating the motion of a human and a robot and also apply it to a robotic wheelchair. They propose to use the velocity obstacles approach [7] to guide the robot to a local virtual target. We compare our method to a similar approach in our experimental section. Most similar to our approach is the work of Morales et al. [12]. They optimize a utility that encodes the desired relative position as well as the walking comfort of the pedestrian. However, they optimize the planned trajectory locally, whereas we optimize future trajectories to a distant subgoal, which allows the robot to adapt its behavior to the environment in a predictive manner.

3 A Socially Compliant Follow the Leader Approach

A better understanding of the natural navigation behavior of pedestrians enables a mobile robot to follow a human leader in a socially more compliant way. In this section, we first formalize the problem of following a leader. We consider the navigation task to stay close to a fixed relative position with respect to its leader. To solve this task, we propose an approach that predicts the trajectory of the pedestrian and at the same time computes a forward-looking trajectory that minimizes the deviation to the desired position.

3.1 Problem Definition

In this work, we consider the 2D navigation behavior of a mobile robot and a leading pedestrian. A trajectory τ_h of the human and τ_r of the robot are mappings $\tau : \mathbb{R} \rightarrow \mathbb{R}^2$ from time to a 2D position. The position of the robot, or the pedestrian, respectively, at time t is thus given by $\tau(t)$ and their velocity by $\dot{\tau}(t)$. We assume a mobile robot with a differential drive that is always oriented in driving direction. Similarly, we assume that the pedestrian is always headed




Fig. 1. The desired position of the robot is a fixed location in the local coordinate system of the pedestrian. The dashed line illustrates the desired trajectory of the robot $\hat{\tau}_r$ given the predicted trajectory $\hat{\tau}_h$ of the pedestrian. Deviation from the desired trajectory yields an additional cost integrated along the trajectory, as illustrated by the shaded area.

in walking direction. Thus, the orientation $\theta(t)$ at time t is the direction of the vector $\dot{\tau}(t)$.

We define the desired position of the robot by a fixed position $\mathbf{o} = (o_x, o_y)$ in the local coordinate system of the pedestrian, i.e., the robot is supposed to always maintain the same position relative to the human. Given the trajectory $\tau_h(t)$ of the human, we can compute the desired trajectory of the robot

$$\hat{\tau}_r(t) = \tau_h(t) + \mathbf{q}(\theta_h(t))\mathbf{o}, \quad (1)$$

where $\mathbf{q}(\theta_h(t))$ is the rotation of the human at time t . In practice the robot cannot always follow this desired trajectory due to obstacles in the environment, or other dynamic constraints. We cast the resulting navigation goal in a utility-optimizing manner, where the cost function is a linear combination of the squared norm of the deviation from the desired trajectory and an additional term $g_{\text{nav}}(\tau, t)$ that comprises acceleration and velocity bounds and clearance to obstacles. Therefore, the desired trajectory minimizes the navigation cost function

$$c(\tau_r) = \int_{t=0}^T (\theta_1 |\tau_r(t) - \hat{\tau}_r(t)|^2 + \theta_2 g_{\text{nav}}(\tau_r, t)) dt, \quad (2)$$

where the weights θ_1 and θ_2 are model parameters to adjust the behavior to the given application. Fig. 1 illustrates the predicted trajectory of the pedestrian, the offset in the local reference frame of the pedestrian and the resulting desired trajectory of the robot. The challenge of this approach is to predict the trajectory of the human, which determines the desired trajectory of the robot $\hat{\tau}_r(t)$. To this end we utilize a predictive model of natural human navigation behavior, which we shortly recap in the following.

3.2 Modeling Human Navigation Behavior

Our approach relies on an accurate model of human navigation behavior that allows the robot to predict the movements of the leading pedestrian. To achieve

socially compliant behavior of the robot, we want to explicitly model the fact that the human is also aware of the robot and reacts to the actions of the robot.

Kretzschmar et al. [9] describe a probabilistic model of such an interactive navigation behavior. For given start and goal positions, the proposed model yields a distribution over the joint space of the trajectories of each agent involved in the navigation process. This probability distribution depends on a weighted sum of features \mathbf{f} that capture important properties of human navigation behavior. Each feature is a function that maps a composite trajectory, i.e., the set of trajectories for all agents, to a real value. Kretzschmar et al. propose features that describe the individual properties of each trajectory, such as the integrated velocity and acceleration along the trajectory, and the time to reach the target. In addition, they propose features that describe interaction between the agents, such as their mutual distance. A weight vector $\boldsymbol{\theta}$ parameterizes the model and describes the importance of each feature in the feature vector \mathbf{f} .

In the special case of two agents h and r , the model yields the distribution

$$p_{\boldsymbol{\theta}}(\boldsymbol{\tau}_h, \boldsymbol{\tau}_r) \propto \exp(-\boldsymbol{\theta}^T \mathbf{f}(\boldsymbol{\tau}_h, \boldsymbol{\tau}_r)), \quad (3)$$

where $\boldsymbol{\tau}_h$ and $\boldsymbol{\tau}_r$ are the trajectories of the two agents, as introduced in the previous section. One can interpret $\boldsymbol{\theta}^T \mathbf{f}(\boldsymbol{\tau}_h, \boldsymbol{\tau}_r)$ as a cost function. The agents are thus exponentially more likely to select a trajectory with lower cost. To adapt the model to the individual navigation behavior of different pedestrians or to a certain environment, we can learn the feature weights $\boldsymbol{\theta}$ from observed data, such that the predicted trajectories accurately resemble the navigation behavior of real humans in the designated environment. Find details on the proposed learning approach as well as a description of features that capture important properties of natural navigation behavior in Kretzschmar et al. [9].

3.3 Unifying Prediction and Planning

We utilize the model proposed by Kretzschmar et al. [9] to predict the trajectory of the pedestrian, and to plan a trajectory for the robot simultaneously. In particular, we adopt the proposed features that capture accelerations, velocities, distances to obstacles and the time to reach the target to predict the natural navigation behavior of the pedestrian. In addition, we introduce the feature

$$f_d(\boldsymbol{\tau}_h, \boldsymbol{\tau}_r) = \int_{t=0}^T |\boldsymbol{\tau}_r(t) - \hat{\boldsymbol{\tau}}_r(t)|^2 dt, \quad (4)$$

that describes the squared deviation from the desired position of the robot along the trajectory, and

$$f_n(\boldsymbol{\tau}_h, \boldsymbol{\tau}_r) = \int_{t=0}^T g_{\text{nav}}(\boldsymbol{\tau}_r, t) dt, \quad (5)$$

to account for further navigation constraints of the robot, as described in Sec. 3.1. During navigation, we compute the most likely composite trajectory $(\boldsymbol{\tau}_h, \boldsymbol{\tau}_r)$ with respect to the probability distribution given by Eq. (3). Due to the additional features f_d and f_n , this most likely composite trajectory not only predicts

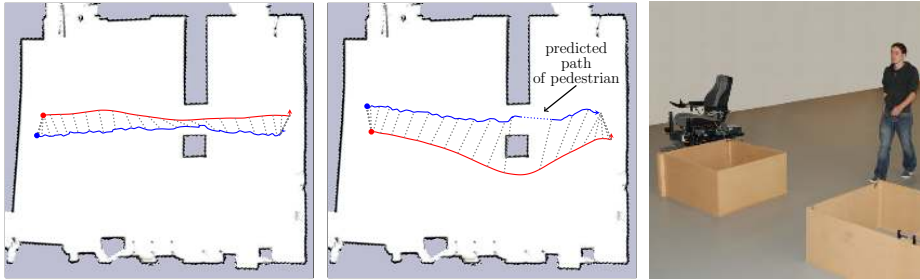


Fig. 2. Left: observed trajectories of the robot (red) and the human (blue) during navigation. The robot falls back behind the pedestrian in the narrow passage. Middle: observed trajectories in an experiment where the robot bypasses the obstacle on the lower side to meet the pedestrian after the passage. Right: Experimental setup.

the trajectory of the pedestrian but also computes the trajectory of the robot that minimizes the navigation cost function of the robot (Eq. (2)). In particular, this method accounts for the effect that the pedestrian interacts with the robot, i.e., that the pedestrian behaves cooperatively and navigates in a way that helps the joint navigation goal. By adjusting the weights of the features we can adapt the level of cooperative behavior that we ascribe to the human. Fig. 3 and Fig. 4 illustrate the predicted trajectory of the pedestrian and the planned trajectory for the robot in two different scenarios.

In addition, the predictive model is beneficial in situations where the leading pedestrian is not in the field of view of the robot’s sensors for some time. Instead of stopping the navigation task, the robot is able to predict the trajectory of the pedestrian and to continue its plan. When the human reappears in the observation of the robot, the people tracker can use the prediction to solve the data association problem, i.e., to select the correct pedestrian as leader.

The predictive model yields predictions of trajectories to known target positions. However, the final target position of the pedestrian is not known in general. In our experiments, we interpolate the observed trajectory of the pedestrian to estimate its target position. In environments where prior information of the typical paths of pedestrians is available, we can also use more sophisticated methods to estimate their target position [2, 18].

4 Experiments

In this section, we describe a set of experiments using a real robotic wheelchair that suggest that our method is applicable to successfully navigate alongside an accompanying pedestrian in the presence of obstacles. Furthermore, we compare our approach in simulation to two related methods. These experiments intend to show the advantages of our predictive planning approach over local control methods, especially in situations where the environmental conditions hinder the robot to remain at its desired position. During the navigation task, our method




Fig. 3. Predictions computed by the wheelchair at four successive time steps. The robot predicts the human to pass the passage. Since the passage is too narrow for the robot (red) and the human (blue) the robot leaves its desired position and let the human pass first. After the passage, the robot resumes its desired position.

continuously computes the most likely composite trajectory by optimizing its probability at a rate of 5 hz.

4.1 Real Robot Experiments

In the following experiments, we use the method proposed in this paper to navigate a robotic wheelchair next to a pedestrian at a distance of 1 m. The robot relies on on-board sensors only. It localizes itself in the environment using Monte Carlo localization [17] and tracks the pedestrian using a laser based people tracker. Fig. 2 shows the paths of the wheelchair and the pedestrian as observed by the wheelchair in two different scenarios.

In the first run (Fig. 2 left), the robot’s desired position is on the left of the pedestrian. It starts moving alongside the pedestrian, falls back behind the pedestrian during passing the passage and catches up afterwards. Fig. 3 shows the predictions of the wheelchair during the navigation task in the same run. As soon as the pedestrian starts to move, the robot computes the most likely composite trajectory of the robot and the pedestrian. It predicts that the pedestrian walks through the passage and that the robot itself stays behind and regains the position to the left of the pedestrian afterwards.

In the second run (Fig. 2 middle), the robot is supposed to keep its position on the right hand side of the pedestrian. Since there is enough space on the lower side of the obstacle, the robot decides to pass the obstacle on a this side, which allows the robot to stay at the human’s side as long as possible. While the pedestrian is in the passage, the obstacle blocks the laser scanner and the robot cannot observe the pedestrian. However, since the robot maintains predictions about the movement of the pedestrian, it is able to follow its planned path and join the pedestrian after it is tracked again. Fig. 4 shows the predictions of the robot during this second run. First, the estimated target is still on the left side of the obstacles due to the low velocity of the pedestrian. However, as soon as the pedestrian proceeds to its goal position, the robot predicts that the pedestrian moves through the passage and plans to pass the obstacle on the other side. While the obstacle occludes the pedestrian, the robot updates its beliefs based on the current prediction of the pedestrian’s position.




Fig. 4. Prediction computed by the wheelchair at four successive time steps. In this experiment, the desired position of the robot is on the right hand side of the human. The robot stays at the human’s side as long as possible. It then evades the obstacle on the right side and continues to move to its desired position relative to the human.

4.2 Comparison in Simulation

Fig. 5 and Fig. 6 show a comparison of our method in simulation to a social forces (SF) based approach [1] and a velocity obstacles (VO) approach, similar to the method proposed by Prassler et al. [14]. To allow for a fair comparison of the methods, we scripted the pedestrian’s path on a rectangular path with a velocity of 0.5 ms^{-1} . The desired position of the robot is 1 m to the left of the pedestrian for all experiments. We set the parameters of all approaches such that the robot always kept a safety distance of at least 0.25 m to the pedestrian, as well as to obstacles in the environment.

Both, SF as well as VO compute control commands towards a virtual target position. To compute this position, we adopt the method proposed by Prassler et al. [14]. They linearly extrapolate the current velocity of the pedestrian in a small time horizon Δt to avoid that the robot lags behind the desired position. We adjust Δt for both methods such that the robot converges to the desired position when the pedestrian moves on a straight line with 0.5 ms^{-1} .

In the test environments, the challenge for the robot is to catch up to the desired position after the pedestrian takes turns on its path. Furthermore, there is a narrow passage in which the robot cannot keep its desired position. Fig. 5 shows that all methods manage to pass the passage. However, the bar plot on the right shows that our method is able to stay closer to the desired position on average along the trajectory. This is due to the fact that our method predicts the trajectory of the pedestrian and computes the trajectory of the robot that minimizes the deviation along the whole path, while also incorporating properties of the robot, such as limited acceleration or velocity constraints. Such long term planning is better suited to accomplish the navigation task compared to greedily approaching the desired position.

Fig. 6 shows a similar experiment with an additional obstacle that resembles an open door in a typical indoor environment. The first image shows that our approach is able to negotiate the passage in a similar way as in the first setup. The robot falls back behind the pedestrian and catches up afterwards. Both SF as well as VO, however, get stuck behind the open door, since there is a local minimum in their local control functions. The bar plot reflects the advantage of the predictive planning in this experiment. Whereas our method shows a similar




Fig. 5. Comparing our method to velocity obstacles (VO) and social forces (SF). The desired position of the robot (red) is one meter to the left side of the human (blue). The bar plot shows that our method stays closer to the desired position on average.




Fig. 6. Comparison to VO and SF that illustrates the advantages of our method over local control methods. Both VO and SF get stuck in the non-convex obstacle.

mean deviation from the desired position as in the first experiment, SF and VO gain a higher deviation whilst stuck in the local minima.

5 Conclusion

In this paper, we presented a novel method that allows a robot to follow a leading person in a socially compliant way. Our approach uses a feature-based model of natural navigation behavior to predict the trajectory of the leading human. In contrast to previous approaches, our method allows the robot to compute far-sighted plans that minimize the long-term deviation from the desired trajectory. In addition to features that describe natural intents of navigating pedestrians, our method uses features that capture the navigation goals of the robot. The resulting model thus unifies prediction of the human’s behavior, and path planning of the robot. In several experiments also carried out with a robotic wheelchair we demonstrated that the proposed model is applicable to real world scenarios such as navigating alongside an accompanying person in the presence of obstacles. A comparison in simulation suggests that our method outperforms previous models that rely on local control strategies.

Acknowledgements

This work has been partially supported by the EC under contract numbers ERC-267686-LifeNav, and FP7-610603-EUROPA2.

Bibliography

- [1] T. Balch and M. Hybinette. Social potentials for scalable multi-robot formations. In *IEEE Int. Conf. on Robotics and Automation*, pages 73–80, 2000.
- [2] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun. Learning motion patterns of people for compliant robot motion. *Int. Journal of Robotics Research*, 24(1): 31–48, 2005.
- [3] Y.-Q. Chen and Z. Wang. Formation control: a review and a new consideration. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 3181–3186, 2005.
- [4] S. Y. Chiem and E. Cervera. Vision-based robot formations with bezier trajectories. In *Intelligent Autonomous Systems*, volume 8, pages 191–198, 2004.
- [5] A. Das, R. Fierro, V. Kumar, J. Ostrowski, J. Spletzer, and C. Taylor. A vision-based formation control framework. *IEEE Transactions on Robotics and Automation*, 18(5):813–825, 2002.
- [6] J. Desai, J. Ostrowski, and V. Kumar. Modeling and control of formations of nonholonomic mobile robots. *IEEE Transactions on Robotics and Automation*, 17(6):905–908, 2001.
- [7] P. Fiorini and Z. Shiller. Motion planning in dynamic environments using velocity obstacles. *Int. Journal of Robotics Research*, 17(7):760–772, 1998.
- [8] J. Huang, S. Farritor, A. Qadi, and S. Goddard. Localization and follow-the-leader control of a heterogeneous group of mobile robots. *IEEE/ASME Transactions on Mechatronics*, 11(2):205–215, 2006.
- [9] H. Kretzschmar, M. Kuderer, and W. Burgard. Learning to predict trajectories of cooperatively navigating agents. In *IEEE Int. Conf. on Robotics and Automation*, 2014.
- [10] S. M. LaValle. *Planning algorithms*. Cambridge University Press, 2006.
- [11] S. Liu, D. Sun, and C. Zhu. Coordinated motion planning of multiple mobile robots in formation. In *World Congress on Intelligent Control and Automation*, pages 1806–1811, 2010.
- [12] Y. Morales, S. Satake, R. Huq, D. Glas, T. Kanda, and N. Hagita. How do people walk side-by-side? using a computational model of human behavior for a social robot. In *ACM/IEEE Int. Conf. on Human-Robot Interaction*, 2012.
- [13] N. Pradhan, T. Burg, S. Birchfield, and U. Hasirci. Indoor navigation for mobile robots using predictive fields. In *American Control Conference*, pages 3237–3241, 2013.
- [14] E. Prassler, D. Bank, B. Kluge, and M. Hagele. Key technologies in robot assistants: Motion coordination between a human and a mobile robot. *Transactions on Control, Automation and Systems Engineering*, 4(1):56–61, 2002.
- [15] L. Qin, Y. Zha, Q. Yin, and Y. Peng. Formation control of robotic swarm using bounded artificial forces. *The Scientific World Journal*, 2013, 2013.
- [16] H. G. Tanner and A. Kumar. Formation stabilization of multiple agents using decentralized navigation functions. In *Robotics: Science and Systems*, pages 49–56, 2005.
- [17] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2):99–141, 2000.
- [18] B. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. Bagnell, M. Hebert, A. Dey, and S. Srinivasa. Planning-based prediction for pedestrians. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 3931–3936, 2009.