

Iterative Human-Aware Mobile Robot Navigation

Shih-Yun Lo, Benito Fernandez, and Peter Stone

University of Texas at Austin

{yunl;pstone}@cs.utexas.edu;benito@utexas.edu

Abstract—Mobile robot navigation in human populated environments has been widely studied in the past two decades. Significant improvements in this technology suggest a promising future for introducing mobile service robots in human workspaces that help with daily activities. State-of-the-art approaches have shown real-world deployments of robots that can safely navigate through dense crowds [14, 12]. Still, most robots lack the ability to navigate in a human-friendly manner, which requires the ability to identify human intentions, especially in collision-risky situations, and to avoid collisions legibly [1]. Most studies to date do not take human-robot interaction into consideration, resulting in unexpected human behaviors during deployment, and ineffective robot planning due to large prediction errors regarding human responses. In this work, we therefore seek to model human reactions around a robot in collision-dangerous scenarios, with specific study of situations in which two agents cross paths. We model people’s collision avoidance behaviors as a function of their underlying intentions and social preferences, and propose to predict human motions based on such intentions and preferences. Preliminary results suggest large variations in different people’s crossing path selection. We turn to a psychological model to explain such diverse responses, and propose to categorize people’s strategies based on their assumptions about the robot.

I. INTRODUCTION

Consider a robot navigating in a busy atrium with people walking towards different destinations. To navigate robustly, the robot must both predict the motions of approaching pedestrians, and plan its own motions so as not to collide with their predicted paths. When doing so, it is not sufficient to consider the paths of pedestrians as fixed. Rather, the robot needs to recognize that its own behaviors can impact those of the pedestrians. That is, the robot and pedestrians are all a part of a tightly coupled multiagent navigation problem.

Many approaches have been proposed in the past decade to simulate human pedestrian dynamics in crowds [4, 8, 6], and have provided insights into human collision avoidance behaviors. However, behaviors observed among human pedestrians may not fully represent their dynamics when interacting with a *robot*. In particular, when deploying a robot in the real world using state-of-the-art approaches, there has been a large gap between the predicted pedestrian motion and observed trajectories [14, 12]. None of the robot planning work in the literature has been based on realistic human-robot dynamics, which has led to inefficient and/or unsafe robot behaviors.

In this paper, we propose to solve the *robot navigation in crowds* problem by taking into account the pedestrian’s reactions to the *robot’s* motions, to ultimately optimize the expected joint efficiency of the robot’s and the pedestrian’s motions. We start with the study of human collision avoidance




Fig. 1. The two pictures illustrate the common patterns of human pedestrians passing the other agent (here the robot) to avoid a potential collision (with the robot moving towards the door, the pedestrian moving to the stair case at the left): (a) to decelerate and wait for the robot passing in front, (b) to accelerate and pass in front of the robot. While slowing down to wait results in shorter paths, speeding up to pass in front is also commonly observed when people navigate in crowds.

behaviors subject to robot’s maneuvers, and pay close attention to scenarios with crossing paths, as illustrated in Fig. 1.

In such situations with anticipated workspace occlusion, pedestrian paths fall into two homotopy classes: either passing in front of the robot (or other pedestrian), or behind it. Such behaviors capture the dynamics of goal-focused people in a workspace occlusion scenario, which can therefore be used to distinguish motions that are directly interacting with the robot, such as intentional blocking.

From gathering real-world human-robot interaction data, we observed that different people have very different strategies over path decisions. We therefore propose a psychological model of their decision making process to capture factors that potentially influence those different avoidance strategies. Such a model helps us identify a particular person’s strategy as it approaches the robot, to further select features to train a predictive model over path decisions.

We also observed that changing the robot behavior can lead to different pedestrian behavior. We therefore propose an iterative learning process by which the robot acts according to its current learned predictive pedestrian motion model to gather revised data to incorporate within a new model. This process is designed to help ensure consistency between the robot’s prediction of pedestrians’ trajectories and their true trajectories in the presence of the mobile robots.

II. RELATED WORK

The earliest attempt to model pedestrian behaviors was through a physics-based modeling approach, in which the spacing between pedestrians follows a potential field to generate collision-safe interactions [4]. Following this work, explicit

models for local collision avoidance were widely studied, to model the subtle adaptive personal-spacing behaviors in pedestrian interactions [11]. Meanwhile, another community makes use of inverse optimal planning techniques [17, 5, 16] to model pedestrian dynamics by learning their policies assuming global observations and rational behaviors. While the above approaches have been shown to predict pedestrian motions offline or in simulation, an issue arose with the Markov Decision Process assumption when the methods were deployed on a robot to navigate in the real world: the real-time dynamics of pedestrians are also dependent on the robot's actions.

In this regard, a multi-player game setting including both human and robot actions in a joint state space was proposed and evaluated on a mobile robot in a fully-observed environment [15]. This game setting has also shown significant results in autonomous cars through the consideration of human actions into their model to achieve communicative robot behaviors [13]. Still, all of the above work are built from data collected by observing human-human interaction, which has degraded planning performance while encountering unmodeled human-robot interactions during real-world deployment [14, 12]. To achieve cooperative interaction with humans and robots sharing a workspace, it is an essential capability for the robot to identify people's intentions, and then plan responsively to their actions. We propose to achieve better human-robot interaction in crowds by learning a predictive model iteratively through online interaction, and propose a predictive model that generates plans to avoid dynamic obstacles based on personal walking features.

III. METHODOLOGY

The following is a receding-horizon control formulation of the constrained optimization problem characterizing the pedestrian's motion. The control $u \in \mathbb{R}^2$ is defined as the acceleration of the pedestrian, with position $x^p \in \mathbb{R}^2$ and velocity $v^p \in \mathbb{R}^2$. We consider discrete system dynamics in continuous state space, so $x_{t+1}^p = x_t^p + v_t^p dt$. Similarly, we denote the "opponent's" position x^o and velocity v^o . We assume that, people plan with a receding-horizon strategy, namely to plan ahead for a fixed horizon from time t to $t+N$, act according to u_t , and start over on the next time step $t+1$:

$$\begin{aligned}
& \min_{u_n \dots u_N} \sum_{t=n}^N C_{length,t} + C_{energy,t} + C_{social,t} + C_{to-go,N} \\
& (r_t^{rel})^2 > (r^{saf})^2 \\
& v_t^p < v_{max} \\
& x_{t+1}^p = x_t^p + v_t^p dt \\
& x_{t+1}^o = x_t^o + v_t^o dt \\
& v_t^p = v_0^p, t \leq n \\
& v_t^o = v_0^o \\
& v_{t+1}^p = v_t^p + u_t dt, t > n.
\end{aligned} \tag{1}$$

We assume that the pedestrian has some rough estimate of his/her own velocity v^p , the opponent's velocity v^o , both agents' desired(initial) velocity v_0^p, v_0^o , and their relative distance $r^{rel} = (x^p - x^o)^T (x^p - x^o)$. The inequality constraints

are posed on the pedestrian's (adaptive) safety margin r^{saf} between two agents, and the pedestrian velocity upper-bound v_{max} . The equality constraints are posed on the dynamics of both agents and their initial velocities v_0^p, v_0^o , which the robot can estimate based on past observations through filtering techniques.

The control input u is not bounded but penalized in C_{energy} for acceleration/velocity variation of the trajectory. C_{length} and C_{to-go} penalize path length and a heuristic cost-to-go. Following a standard quadratic formulation in optimal control, we have: $C_{energy,t} = u_t^T R u_t$, $C_{length,t} = \delta_t^T Q \delta_t$, $\delta_t = x_{t+1} - x_t$, $C_{N,to-go} = (x_N - x^G)^T P_N (x_N - x^G)$, where x^G is the (sub)goal position of the pedestrian. P_N can be solved through dynamic programming: $P_N = Q + K_N^T R K_N + (A + B K_N)^T P_N (A + B K_N)$, $K_N = -(R + B^T P_{N-1} B)^{-1} B^T P_{N-1} A$, considering a linear system dynamics $s_{t+1}^p = A s_t^p + B u_t$, here $s^p = [x^p, v^p]$. Q and R are positive semi-definite matrices. Other forms of cost-to-go can also be adopted, such as naive distance estimate (considering only distance-to-go), or social cost-to-go estimate (such as considering the distance of x_N^p to a large group of static people).

C_{social} represents the pedestrian's social interaction tendencies when navigating in a crowd. More specifically, based on studies of robot legible motion in navigation [7, 10] and human spacing/dynamics in crowds [4, 2], we design the social cost function C_{social} to include *visibility*, and *comfort*, to explain human crossing path selection:

$$\begin{aligned}
C_{social} &= \omega^T f(x), \\
f(x) &= [C_{vis}, C_{comf}],
\end{aligned} \tag{2}$$

where ω is a weighting function. As suggested in [7], when planning for crossing, people prefer the traversing path that maintains visibility from the point of view of the other agent (i.e. passing in front). From the social force formulation [4], velocity with a large component along the other agent's opposite direction causes higher repellent forces. We model this repellent force to be caused by a cost function referred as social discomfort, which appears when the opponent agent is within the pedestrian's social space:

$$C_{comf} = (v^p - v^o)^T (v^p) \mathbb{1}(r^{rel} < r^{soc}). \tag{3}$$

This term penalizes crossing behaviors that turn towards the opponent agent. Along with visibility assessment, the pedestrian who chooses to pass behind has an even higher cost to speed up compared with the maneuver to pass in front. This contributes to the commonly seen slowing-down behavior in a^p , which results in the pedestrian staying mostly on the original route, shown in Fig. 1(a).

Note that we build into our model a (adaptive) response time n before the person acts according to the cost function, following the notion of *psychological tension* before reacting to mental decisions [4]. This variable, reflected in Eq. 1 allows us to capture the fact that the person may have decided to change paths before he/she actually takes the action.

The crossing path decisions vary much among different people, which we can explain through the variables in Eq. 1,

such as n (response time), r^{saf} (safety margin), ω (social cost weighting function) and possibly the weighting on other cost terms in Eq. 1. Those psychological and social preference factors of a person can also vary based on the type of "opponent", for example, a robot rather than a human.

As suggested in [3], people have different degrees of interest/comfort in interacting with a robot, which can affect their maneuvers during human-robot crossing. For example, in human crowd dynamics simulation, the safety margin is considered merely a function of crowd density and walking speed [4], whereas in real-world robot deployments, we observe that this variable can vary a lot from person to person.

To predict human-robot crossing behaviors, it is therefore insufficient to learn one set of parameters from one large set of interaction data across many people. Further, the change in the pedestrian's dynamic behavior can be sudden when he/she reacts to his/her crossing decisions (considering pedestrian dynamics change at time n in Eq. 1), and the trajectories associated with the different crossing decisions, such as to pass from behind or in front, may have very different dynamical performance, i.e. maximum speed and acceleration patterns. Therefore, inaccurate crossing path prediction can negatively affect the joint path efficiency, and possibly violate the safety criteria.

However, to identify appropriate values for all the variables that affect the pedestrian's motion (i.e. values that describe their observed motion so far) is computationally intractable from a model training perspective. For example, for inverse optimal planning, the algorithm seeks to find a fixed ω in Eq. 2 that maximizes the likelihood function P of observed trajectories ξ ,

$$\max_{\omega} P(\xi|\omega), \quad (4)$$

and trajectories representing different values of ω need to first be classified, or the algorithm will not converge.

Still, to plan in the crowd, humans do a successful job without accurate prediction, and false prediction is even commonly observed without necessary safety harm. For example, with a direct confronting agent in the front walking towards you, people sometimes avoid to the same side at the same time (or within close time). What seems to matter is then just *which way* to go instead of the entire trajectory prediction.

Even so, when moving through crowds, humans succeed at avoiding collisions without fully accurate predictions about other people's motion. For example, when two pedestrians directly approach one another, they sometimes move to the same side at the same time (or similar time). When that happens, they quickly respond to resolve the conflict. What seems to matter is just *which way* to go, rather than the entire trajectory prediction.

We therefore propose to seek for a much simpler predictor over pedestrian path crossing decision manifold: to pass from the front (a^P) or behind (a^A), through logistic regression,

$$p(a = a^P | \xi_t^P, s_t) \sim \frac{1}{1 + e^{w^T f(\xi_t^P, s_t)}}. \quad (5)$$

Here, $s_t = [x_t^P, v_t^P, x_t^o, v_t^o]$ is the state of the overall system, included to describe the crossing scenario; ξ_t^P is the past trajectory of the pedestrian at time t , included to describe the pedestrian's dynamics; f is a feature function, and w is a weighting to be learnt.

To identify the crossing situation, f is chosen to present features such as deviations from the desired heading of both agents, their average velocities, relative distance, and the angle between robot position and pedestrian goal, from the point of view of the pedestrian.

The features we choose to describe the pedestrian's dynamics include: his/her average velocity, velocity variance, and heading variance.

Observations of past interactions are also informative of the person's path crossing preferences. For example, we can count the number of passive crossing p and active crossing q , and pose a Beta distribution as the conjugate prior over the passive crossing decision a^P as: $Beta(p+1, q+1)$. However, diagnosis of past interactions is not necessarily easy to come by with limited sensing.

We propose to iteratively learn the parameter w in this predictive model by collecting data from real-world human-robot interactions. Since people are also adaptive about their reaction to a robot, new types of behaviors may be observed as the robot changes its behavior.

A. Intention Expression and Opponent Response Hypothesis during Human Crossing

When a pedestrian decides the path for crossing, the *dynamic patterns* of his/her motion, such as speeding up at the same time turning away, indicate the pedestrian's intended path decision, which we refer in this paper as the *intention*: the intention to actively cross the other agent from the front by speeding up (a^A), or to passively cross from behind by slowing down (a^P).

When a pedestrian decides what path to take in order to cross the robot's path, the early part of his or her trajectory, such as speeding up or turning away, are indicative of the pedestrian's path *intention*: the intention to actively cross the other agent from the front by speeding up (a^A), or to passively cross from behind by slowing down (a^P).

A robot's predictability and intention expressiveness have been identified as important criteria for designing social-friendly or cooperative robots in the human-robot interaction literature. In our human-robot crossing scenario, it is then especially important for the person to be able to *predict* the robot's motion, so as to feel *safe* around a robot. We propose to achieve predictable and intention-expressive motion by mimicking the dynamic patterns of people's crossing behaviors, so as to communicate the robot's intention of crossing in a human-familiar manner.

Intention communication is commonly observed in crowd navigation, to allow two self-interested agents to reach consensus on the ways they are going to avoid one another, based on their indicated future trajectories. To enable the robot to communicate with a human pedestrian and to predict the

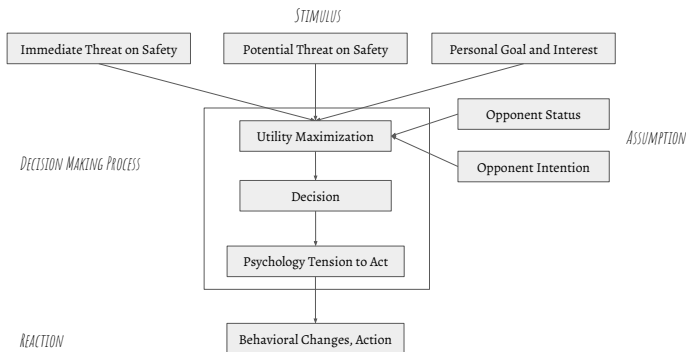


Fig. 2. We propose this schematic illustration of human action model based on the psychological process analysis on behavioral changes in [4]. Here, we further distinguish stimulus, and propose to include opponent behavior into the decision making process.

resultant behavior change of the confronting pedestrian, we propose a predictive model following the formulation in Eq. 5, taking into consideration the robot’s action a^r (designed as to follow human crossing dynamic patterns):

$$p(a = a^P | \xi_t^P, s_t, a^r) \sim \frac{1}{1 + e^{w^T f(\xi_t^P, s_t, a^r)}}. \quad (6)$$

Note that, enabling predictable motions not only enhances human trust in the robot, but also reduces the noise in the crossing pedestrian’s motion that is caused by his/her uncertainty of the robot’s future behavior.

Such a model enables the robot to predict the pedestrian’s motion, based on the robot’s crossing action, and facilitates realistic look-ahead planning by the robot. Similarly to training the free variables in Eq. 5, we propose to iteratively learn this predictive model from real-world robot deployments. Since people also adapt their reactions to a robot, new types of behaviors may be observed over time, as the robot changes its behavior.

B. Human Decision Model considering Opponent Status and Intention Hypothesis

When running the experiments described below in Section III-C, people commonly described themselves as “aware” of the robot, which we interpreted as “requiring attention”. Those people also suggested that they were not navigating as comfortably as they did in human crowds. However, some other people had the opposite opinion, suggesting that the same robot is comfortable to navigate around.

Going back to Eq. 1, we desire to model the pedestrian’s crossing response over a robot, and we desire to understand the mechanism behind the decision making process over crossing behaviors, so we can answer: why do the responses differ so much among different people? We start by proposing a hypothesis that people have very different perception of the robot’s behavior and intention, and introduce a psychological model of human decision making process over crossing strategies, as shown in Fig. 2.

This model takes into account their personal goal and environmental variables, and their hypothesis over the opponent

behavior, here, the robot. We consider sources of stimulus (the input of the pedestrian’s decision making process) with different levels of response time (psychological tension) and frequency of that instance to vary. For example, goal and interest are usually fixed during the entire navigation process. Immediate instances such as unexpected turns of nearby agents cause direct threats and therefore stimulate quick reaction.

As for our human crossing scenario with a robot, we model people as having different assumptions regarding the following items, and use the model to describe the behavior observed in our experiments described in Section III-C:

- 1) the perception of robot identity: is the robot non-self-interested or self-interested?
- 2) the capability to maintain safety: is the robot safe to be around? Can the robot see me from here?
- 3) the capability to communicate intention: is the robot going to understand my intention? Will the robot give way if I walk in front of it?

C. Preliminary Results

We conduct experiments in an atrium with the robot following pre-assigned routes. We first show the participants how the robot runs the task, and then ask them to travel between four sets of waypoints such that collisions with the robot are possible. Each participant does the experiment twice, with different values of safety margin (0.3m and 1m) that the robot keeps from the pedestrian. After the tests, we asked people how they felt about navigating around the robot. Six people participated in the test, none of whom had any prior experience with the robot. We adopt a ROS people-tracking package from [9] to record pedestrian trajectories with an onboard Velodyne laser scanner.

Based on the observations on how far people kept from the robot (the safety margins), and how they interacted with the robot during crossing, we informally categorize people into four groups, displaying different assumptions about the robot:

- 1) Cautious people, who assume the robot to be unsafe and avoid the robot,
- 2) Overly-comfortable people, who assume the robot to be safe and not self-interested. They assume the robot will always give way to human;
- 3) Comfortable, interactive people, who assume the robot to be safe and are willing to indicate intention. They change strategies when the robot’s action is mismatched with their expectation;
- 4) Aggressive people, who largely prefer to actively pass the robot, possibly resulting in large deviation from their original paths.

In our ongoing research, we intend to collect more data to allow our robot to robustly predict human path crossing decisions before and after the robot indicates its intention (shown in Eq. 5 6), following the iterative retraining methodology proposed in this paper.

REFERENCES

- [1] Anca D Dragan, Kenton CT Lee, and Siddhartha S Srinivasa. Legibility and predictability of robot motion. In *Human-Robot Interaction (HRI)*,

- 2013 8th ACM/IEEE International Conference on, pages 301–308. IEEE, 2013.
- [2] Edward Twitchell Hall. The hidden dimension. 1966.
 - [3] Soren Tranberg Hansen, Mikael Svenstrup, Hans Jorgen Andersen, and Thomas Bak. Adaptive human aware navigation based on motion pattern analysis. In *Robot and Human Interactive Communication, 2009. RO-MAN 2009. The 18th IEEE International Symposium on*, pages 927–932. IEEE, 2009.
 - [4] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.
 - [5] Peter Henry, Christian Vollmer, Brian Ferris, and Dieter Fox. Learning to navigate through crowded environments. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 981–986. IEEE, 2010.
 - [6] Ioannis Karamouzas, Peter Heil, Pascal van Beek, and Mark H Overmars. A predictive collision avoidance model for pedestrian simulation. In *International Workshop on Motion in Games*, pages 41–52. Springer, 2009.
 - [7] Thibault Kruse, Patrizia Basili, Stefan Glasauer, and Alexandra Kirsch. Legible robot navigation in the proximity of moving humans. In *Advanced Robotics and its Social Impacts (ARSO), 2012 IEEE Workshop on*, pages 83–88. IEEE, 2012.
 - [8] Fabrice Lamarche and Stéphane Donikian. Crowd of virtual humans: a new approach for real time navigation in complex and structured environments. In *Computer Graphics Forum*, volume 23, pages 509–518. Wiley Online Library, 2004.
 - [9] Angus Leigh, Joelle Pineau, Nicolas Olmedo, and Hong Zhang. Person tracking and following with 2d laser scanners. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 726–733. IEEE, 2015.
 - [10] Christina Lichtenthäler, Tamara Lorenzy, and Alexandra Kirsch. Influence of legibility on perceived safety in a virtual human-robot path crossing task. In *RO-MAN, 2012 IEEE*, pages 676–681. IEEE, 2012.
 - [11] Panagiotis Papadakis, Patrick Rives, and Anne Spalanzani. Adaptive spacing in human-robot interactions. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2627–2632. IEEE, 2014.
 - [12] Mark Pfeiffer, Ulrich Schwesinger, Hannes Sommer, Enric Galceran, and Roland Siegwart. Predicting actions to act predictably: Cooperative partial motion planning with maximum entropy models. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 2096–2101. IEEE, 2016.
 - [13] Dorsa Sadigh, Shankar Sastry, Sanjit A Seshia, and Anca D Dragan. Planning for autonomous cars that leverages effects on human actions. In *Proceedings of the Robotics: Science and Systems Conference (RSS)*, 2016.
 - [14] Pete Trautman, Jeremy Ma, Richard M Murray, and Andreas Krause. Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation. *The International Journal of Robotics Research*, 34(3):335–356, 2015.
 - [15] Peter Trautman and Andreas Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 797–803. IEEE, 2010.
 - [16] Dizan Vasquez, Billy Okal, and Kai O Arras. Inverse reinforcement learning algorithms and features for robot navigation in crowds: an experimental comparison. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1341–1346. IEEE, 2014.
 - [17] Brian D Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, J Andrew Bagnell, Martial Hebert, Anind K Dey, and Siddhartha Srinivasa. Planning-based prediction for pedestrians. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 3931–3936. IEEE, 2009.