

Look Before You Leap: Socially Acceptable High-Speed Ground Robot Navigation in Crowded Hallways

Lakshay Sharma¹, Nicolaniello Buono¹, Ashton Flather¹, Xiaoyi Cai¹, and Jonathan P. How¹

Abstract—To operate safely and efficiently, autonomous warehouse/delivery robots must be able to accomplish tasks while navigating in dynamic environments and handling the large uncertainties associated with the motions/behaviors of other robots and/or humans. A key scenario in such environments is the hallway problem, where robots must operate in the same narrow corridor as human traffic going in one or both directions. Traditionally, robot planners have tended to focus on socially acceptable behavior in the hallway scenario at the expense of performance. This paper proposes a planner that aims to address the consequent “robot freezing problem” in hallways by allowing for “peek-and-pass” maneuvers. We then go on to demonstrate in simulation how this planner improves robot time to goal without violating social norms. Finally, we show initial hardware demonstrations of this planner in the real world, along with a novel STAR (Socially Trained Agile Robot) platform designed with human comfort in mind.

I. INTRODUCTION

Human-robot co-existence and collaboration is essential to expanding current automation capabilities. Social robot navigation, which comprises as planning robot behavior in a dynamic, uncertain environment co-occupied by humans is therefore of increasing importance. Increasing robot and human productivity in a shared workspace requires improvements in two key areas:

- Average robot speed: improvements are needed to maximize robot productivity and minimize inter-stage latency
- Human comfort with robot behavior: with humans feeling uneasy or unsafe around robots in a professional environment, productivity of both agents drops

This leads to the realization that there is a need for high speed, safe and socially mobile robot motion planners that can safely negotiate space while obeying appropriate social politeness norms. Now, it is usually noted that the two objectives above demonstrate trade-offs against each other. Additionally, human comfort is a subjective metric which is hard to quantify accurately, leading to robots usually being isolated to prevent human injury in a workspace. Consequently, it is all too common for state of the art social navigation algorithms to pick one extreme or another: either planners are too aggressive, leading to rude and/or unpredictable behaviors which make nearby humans uneasy (less common), or planners are too conservative, leading to robots being stuck behind slow pedestrians and unable to complete their tasks efficiently (more common).

¹Aerospace Controls Laboratory, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139, USA {lakshays, nbuono, ashton.f, xyc, jhow}@mit.edu

This paper therefore attempts to find a middle ground, by proposing a socially acceptable planner that solves the “robot freezing problem” of being stuck behind a slow human, by executing an overtaking maneuver. However, to do this safely, the robot must be careful to look before it leaps: peeking before passing to ensure safety for itself and any unseen human who might be occupying the unobserved space needed to pass.

II. RELATED WORK

Social navigation is a rich field of study, and there have been many attempts at solving these issues. We can broadly categorize planners into two kinds: learning-based and optimization-based.

Learning-based planners such as [1] are rapidly emerging as a robust way to plan in social environments. However, we choose to focus on optimization-based planners such as [2], [3] for our approach to solve the social navigation issues under consideration. This is because safety and robustness guarantees for learning-based planning pipelines are not fully resolved yet, and optimization-based planners remain deterministically safe especially in safety-critical scenarios.

There is a lot of research into the “freezing robot problem”, which arises from the physical inability of a robot to move through dense crowds without accounting for human cooperation. [4], [5], [6] discuss the freezing robot problem and the reciprocal dance, which are the key challenges we aim to address in this paper.

There has also been much work in the field of risk metrics for robots. In particular, [7] and [8] discuss risk metrics that should be used in robotics. In this paper, we use CVaR as described in [7] as a framework to judge the risk of collision given the current belief from sensor data, and use it to incentivize peeking that enables the de-risking of high-uncertainty regions by active sensing when necessary.

Recent work done in [9] mentions several key metrics using which human comfort can be measured, which are rapidly being adopted by the community. We use several of these metrics to quantify different versions of our planner against the chosen baselines. In particular, with time to goal as the key performance indicator, we use number of human path changes and number of times the robot enters a human’s personal space as measures of “rudeness”. There are several variants of human-robot distance as a metric supported and used by existing work, such as [10], [11], [12]. In particular, [12] discusses the idea of social transit as an important feature to maximize safety and comfort when people share environments with social robots. Human

trajectory adjustment frequency is a less commonly used metric that we think is valuable in the realistic scenario of smart humans avoiding collisions with a robot before they happen.

III. PIPELINE OVERVIEW

Our pipeline, outlined in Fig. 1 consists of three levels: mapping (conversion of sensor data to a 2D costmap), prediction (predicting future human trajectories for dynamic obstacle avoidance), and planning (generation and execution of a motion plan given the costmap).

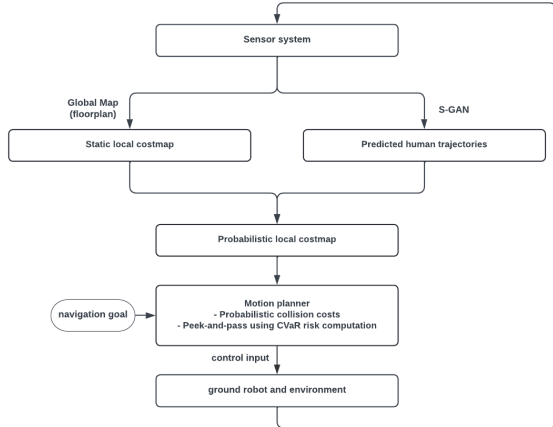


Fig. 1. Overview of the proposed planning pipeline.

A. Mapping

The robot is assumed to have a global map of its environment so that walls, shelves and other permanent static obstacles are predetermined. Additionally, we take as given the following sensors: a 2D laser scanner, and two stereo depth cameras – one pointing forwards and one pointing backwards. These sensors are then used to generate an occupancy map, with obstacle inflation then producing a 2D costmap for the instantaneous environment state.

B. Prediction

The 2D costmap generated above is not sufficient to plan in an environment with dynamic human agents. Therefore, we generate future human trajectories using an S-GAN based online predictor, similar to [13]. This creates a time-dependent distribution for occupancy probabilities, and therefore costs, that is overlaid on the costmap to be used by the planner.

C. Planning

Given the 2D costmap with static obstacles and distributions of positions, costs and occupancy probabilities for dynamic obstacles, this paper’s key contribution is a novel risk-aware motion primitives-based robot motion planner that targets the robot freezing problem in hallways. By allowing for peek-and-pass maneuvers, we demonstrate improved times to goal without sacrificing social acceptability. The details of this planner are described below in Section IV.

IV. PLANNER

With the costmap described in Sections III A and B above, we can proceed with generating a robust motion plan for the robot to execute. We consider here the hallway passing problem, described in Fig. 2.

In a social robot navigation scenario, especially in a spatially narrow setting such as a hallway, human comfort with robot motion is critical. One of the most common ways that robots can induce social discomfort is by executing jerky motions and sharp turns that greatly reduce robot predictability. Therefore, we opt to use a motion primitives based approach similar to [14], since by generating plans that are guaranteed to respect vehicle dynamics they offer greater social comfort at high speeds than MPC-based methods.

Aside from social acceptability, the key risk that we must account for when going fast in a hallway is that of humans approaching from the opposite direction in the unknown space created by a human right in front of the robot. To achieve this, we use the CVaR metric [7] to predict the risk of future robot plans into unknown space.

A. Problem Formulation

We work with a differential drive ground robot operating in a 2D plane, subject to unicycle dynamics:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} \cos \theta & 0 \\ \sin \theta & 0 \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} v \\ \omega \end{bmatrix} \quad (1)$$

where x and y are the coordinates of the robot in the $x - y$ plane it is operating in, and θ is the orientation of the robot with respect to a designated front.

Now consider the hallway passing scenario shown in Fig. 2. We first formulate an optimization problem, which when solved gives us a plan with the following properties:

- High average speed (low time to goal) for the robot
- Obstacle avoidance, specifically:
 - deterministic obstacle avoidance for static obstacles, and
 - stochastic risk-based avoidance for dynamic obstacles (humans)
- Cautious "peek-and-go" behavior when the robot is blocked by a slow human
- Social acceptability of robot motion, particularly minimizing:
 - personal space violations, and
 - human trajectory interruptions (number of times a human has to change their path due to the robot cutting them off)

To get these properties, the optimization problem to be solved is:

$$\mathbf{x}^*(t) = \arg \min_{\mathbf{x}} J(\mathbf{x}(t), \mathbf{u}) \quad (2)$$

where $J(\mathbf{x}(t), \mathbf{u})$ is a cost function that embodies the desired properties of goal-seeking, safe and socially acceptable exploration behavior outlined above.




Fig. 2. Hallway passing scenario for a fast ground robot stuck behind a slow human.

B. Cost Function

Referencing Fig. 2, for a candidate trajectory $\mathbf{x}(t)$ obtained as a sequence of motion primitives, we then formulate a cost function for the planner to optimize as follows:

$$J(\mathbf{x}(t), \mathbf{u}) = c_1 \|\mathbf{x}(t_f) - \mathbf{x}_g\|^2 + c_2 p_{\text{coll}}(\mathbf{x}(t)) + c_3 \rho(\mathbf{x}(t)), \quad (3)$$

$$\text{s.t. } \dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u}, \quad (4)$$

$$p_{\text{coll}}(\mathbf{x}(t)) = \sum_{i=0}^k p_{\text{coll}}^i(\mathbf{x}(t)), \quad (5)$$

$$\rho(\mathbf{x}(t)) = \max(0, \text{CVaR}_{\alpha_{\text{peek}}}(\Psi(x_k(t), Y(t)) - \epsilon_{\text{peek}}), \quad (6)$$

where c_1, c_2 and c_3 are tunable scalar cost parameters, α_{peek} and ϵ_{peek} are tunable risk parameters for the peeking maneuver, t_f is the planning horizon of the planner, \mathbf{x}_g is the goal state, k is the number of humans/dynamic agents being tracked and predicted, $p_{\text{coll}}(\mathbf{x}(t))$ and $\rho(\mathbf{x}(t))$ are the risks of collision for a candidate trajectory $\mathbf{x}(t)$ for observable and unobservable obstacles, respectively. Note that $p_{\text{coll}}^0(\mathbf{x}(t))$ is the chance of collision with a static observable obstacle, calculable directly from the 2D occupancy map, whereas $p_{\text{coll}}^i(\mathbf{x}(t)) \forall i \neq 0$ is the chance of collision with the i^{th} observed human/dynamic agent, calculated as a CVaR cost from the predicted trajectory distribution.




Fig. 3. A laned simulation environment where a robot (green) navigates safely around several humans (blue) which are all moving left to right at varying speeds.

Focusing more on $\rho(\mathbf{x}(t))$, we first formulate a distribution Ψ representing the possible trajectories a human moving from an unobserved region of the map into the currently observable local map (denoted here as $Y(t)$). This distribution, sampled uniformly, is responsible for the emergence of peeking behavior rather than full-speed overtaking in the resultant plans. This is because a full-speed passing maneuver incurs two kinds of risk:

- risk of collision with an unseen human on the passing side, and
- risk of collision with an unseen human in front of the human being passed.

Since the distribution Ψ captures both of these risks and $\rho(\mathbf{x}(t))$ penalizes them accordingly, our planner incentivizes peeking before passing, rather than risky overtaking at high speeds without sufficient information.

V. RESULTS

The motion planning pipeline described thus far is tested against several other planners in simulation, and is also demonstrated on a real robot in a hallway.

A. Preliminary Laned Simulation Results

To first demonstrate that our planner improves performance over a conservative, non-peeking approach, we use a laned simulation environment called HighwayEnv [15] shown in Fig. 3. This allows us to test the planner's performance in a well-structured environment.

Using this environment to test a robot with a maximum speed of 3 m/s, we see in Fig. 4 that as the exploration to exploitation parameter ratio for the planner (defined as $\frac{1}{c_1 c_3}$ using the notation in Eq. 3) increases, higher average speeds are achieved, until an optimal value, beyond which the robot loses a strong sense of attraction for the goal. In particular:

- Low exploration rates lead to low average speeds for the robot, with wide distributions of behavior that reflect reliance on lucky gaps
- High exploration rates show diminishing returns and even performance reduction, as the robot starts to emphasize uncertainty reduction over progress to the goal
- Optimal behavior is shown as a significant increase in speed at intermediate values of $\frac{1}{c_1 c_3}$ (exploration vs. exploitation)

These results show that the proposed idea of exploration/"peeking to pass" successfully increases robot performance without sacrificing safety. We also use this environment to tune the cost parameters.

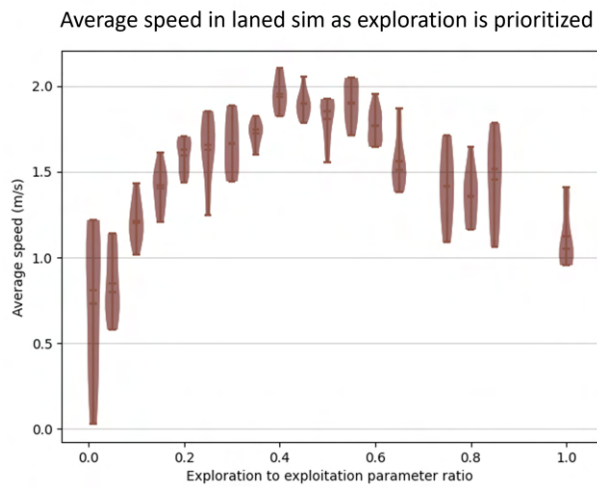


Fig. 4. The laned simulation environment demonstrates that in a structured scenario, the proposed idea of “peek-and-go” successfully increases performance by actively sensing and clearing unknown space.




Fig. 5. Unlaned Environment

B. Unlaned Photorealistic Simulation Results

Having verified the performance gains of our approach above, we move to a photorealistic unlaned simulation in order to show detailed comparisons against baselines. We used NVIDIA Isaac Sim to recreate a crowded hallway environment with people walking both ways, as shown in Fig. 5. The length of this hallway was varied between 25 m, 50 m and 100 m, and the allowable human speeds (fixed during a single trial) range from 0 m/s to 3.5 m/s in increments of 0.25 m/s. A differential drive robot with a top speed of 3 m/s is spawned into this environment near one end of the hallway for each trial, and given a goal at the other end of the hallway.

The proposed pipeline is compared to three other planners: CADRL [1], PSMM [2] and HATEB [3]. The results are shown in Fig. 6, indicate that our planner outperforms commonly used state of the art planners in this uncertain social navigation scenario, with consistent results that are scalable over long distances and across a variety of human speeds.

In particular, we make the following observations about performance (times to goal):

- In short corridors of length 25 m, our planner outperforms the baselines considered, with the average speeds showing a reasonable improvement. However, due to the short length of the corridors, the times to goal for all the planners are close and within each other’s error bars.
- On increasing the lengths of the hallway to 50 m, the speed advantage that our planner has is maintained, and the gap between the baselines increases.
- On increasing corridor length to 100 m, the advantage that our planner possesses is clear to see, and we observe improved performance, save for a few outliers.

Next, we compare the proposed planner to the three baselines for two important metrics of social acceptability of robot motion: the number of times a human has to change their path to avoid the robot, and the number of times a robot enters a human’s personal space (defined here as 0.5 m). We demonstrate that despite the increase in performance, our planner does not appreciably sacrifice social acceptability over more conservative options. In particular, we observe that:

- Our planner shows a greatly reduced number of human path changes compared to the baselines considered. In particular, HATEB is seen to be comparable to our method, with the proposed planner giving generally lower numbers of path changes than HATEB except for a few outliers. These cases are generally seen to arise due to the robot having to merge back in after a failed peek, slowing down humans behind it.
- The proposed pipeline is slightly worse when compared to HATEB when considering the number of times it enters a human’s personal space. However, this can be attributed to our planner having to violate personal space bubbles when passing a slow human, and it is noted that such proximity usually does not make humans uncomfortable. It is therefore key to consider human path changes occurring due to the robot in tandem with personal space violations, in order to get a bigger picture.

Lastly, we analyze how robot performance varies in a 50 m long corridor when human speeds are changed from 0 m/s to 2.5 m/s (a reasonable upper limit for a brisk walk by an average human). We also compare our planner’s behavior to itself when peeking is turned on or off (by making c_3 optimal or high). We make the following observations:

- With peeking enabled, when human speed increases:
 - Initially, the robot travels at close to its maximum speed, with very slow or stationary humans simply being static obstacles that are easy to avoid.
 - Beyond the static regime, the robot is still able to maintain near-optimal performance, due to the peeking and passing maneuver that our planner allows.
 - Increasing human prediction uncertainty with increasing human speed beyond approximately 1.25 m/s starts to force the robot to follow the human in front of it instead of attempting a peek-

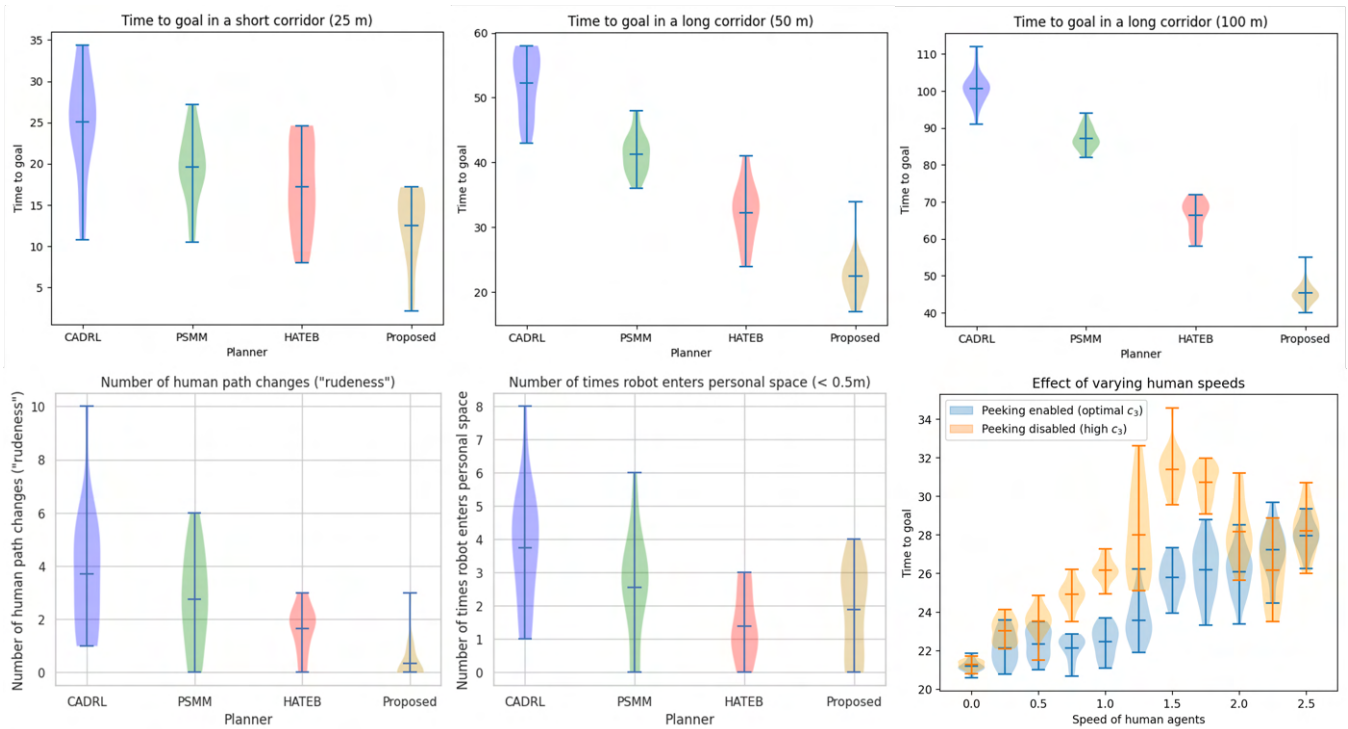


Fig. 6. Trends seen when comparing the proposed planner to the selected baselines. The proposed planner outperforms commonly used state of the art planners in this uncertain social navigation scenario, with consistent results that are scalable over long distances and across a variety of human speeds.

ing/passing maneuver, thereby slowing it down.

- Further increasing the human speed results in a “rear-end panic”, where the predicted trajectories of fast humans behind the robot lead it to believe that a collision is imminent, thereby creating unnecessary “dodging” maneuvers.
- With peeking disabled, when human speed increases:
 - Initially, even at low human speeds, the robot is more hesitant to go around slow humans due to fear of unknown space.
 - As human speed increases to roughly 1.25 m/s, the predicted trajectory spread for a human in front of the robot begins to occupy the entire width of the corridor, with a high enough probability to block the risk-averse robot. This is when the robot switches from going around humans to following humans, and we see a sharp decline in performance (increase in time to goal).
 - Beyond 1.25 m/s, the robot is simply following a faster human, so we observe increasing performance. (with the caveat of “rear-end panic” still present).

These results therefore show that our proposed planner holds promise as a socially acceptable solution for high-speed robot planning in crowded, dynamic environments.

C. Hardware Demonstration

The pipeline described thus far was run on differential drive ground robot platforms, in two hallway scenarios


with different crowd densities. Fig. 7(a) shows a novel STAR (Socially Trained Agile Robot) platform designed with human comfort in mind, with features such as spring-loaded bumpers, covered wheels, low profile and very low operating noise. Fig. 7 shows another robot navigating the less crowded scenario. In both scenarios, it was observed that the planner was able to follow humans in front of it at an acceptable distance without causing discomfort. Additionally, in the first scenario, the robot was seen to slow down and eventually pass a human in front of it who exhibits unexpected slowdowns. It was also observed that the robot postpones the passing maneuver if it is too risky, such as in Fig. 7 when there is an unobserved human coming down the passing lane.

VI. CONCLUSIONS AND FUTURE WORK


This paper presented an effective robot motion planner to navigate crowded and dynamic hallway scenarios at high speeds, by allowing a robot to execute “peek-and-pass” maneuvers. We demonstrated how passing a slow pedestrian can lead to significant performance gains, and offered a way to solve the “robot freezing problem”. This was supported by multiple detailed simulation experiments and hardware demonstrations.

For future work, the authors plan to do much more extensive testing with a faster robot that can reach a higher top speed of 3 m/s or more. This is essential due to

Additionally, the “rear-end panic” issue described at the end of Section V-B, while not likely to occur in a scenario where humans are moving at reasonable speeds, is



(a)



(b)

Fig. 7. (a) The novel STAR (Socially Trained Agile Robot) designed with human comfort in mind. (b) The proposed pipeline was shown to successfully avoid pedestrians and pass slow humans, while canceling/postponing the passing maneuver if the initial peek deems it to be too risky (such as in the case of an oncoming pedestrian) .

nevertheless still an obstacle to effective and socially acceptable navigation. This is compounded by the factor that in a scenario where such decisions do matter, such as a hospital, the stakes are much higher and this property of the planner is not allowable. Addressing robot reaction to human trajectories colliding with it is a problem of two-way human-robot interactions, and we plan to draw from the rich body of work that deals with this issue.

Lastly, high speed planning entails low computation time, which has long plagued optimization-based planners such as the one presented here. We plan to look into imitation learning as an effective and proven way to accelerate our planner to handle higher speeds on smaller platforms.

ACKNOWLEDGMENTS

This work was funded by Amazon under the Amazon Science Hub collaboration with MIT. The support, insightful comments and video help from Zhiye Song, Kota Kondo, Jeremy Cai, Dr. Andrea Tagliabue, Aneesa Sonawalla, Mia Tian, Naman Aggarwal and Ne Myo Han are greatly appreciated.

REFERENCES

- [1] Y. F. Chen, M. Liu, M. Everett, and J. P. How, “Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning,” in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 285–292.
- [2] X.-T. Truong and T. D. Ngo, “Toward socially aware robot navigation in dynamic and crowded environments: A proactive social motion model,” *IEEE Transactions on Automation Science and Engineering*, vol. 14, no. 4, pp. 1743–1760, 2017.
- [3] P. T. Singamaneni, A. Favier, and R. Alami, “Human-aware navigation planner for diverse human-robot interaction contexts,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5817–5824.
- [4] P. Trautman and A. Krause, “Unfreezing the robot: Navigation in dense, interacting crowds,” in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 797–803.
- [5] P. Trautman, J. Ma, R. M. Murray, and A. Krause, “Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation,” *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 335–356, 2015.
- [6] F. Feurtey, “Simulating the collision avoidance behavior of pedestrians,” *The University of Tokyo, School of Engineering, Department of Electronic Engineering*, 2000.
- [7] A. Majumdar and M. Pavone, “How should a robot assess risk? towards an axiomatic theory of risk in robotics,” in *Robotics Research: The 18th International Symposium ISRR*. Springer, 2020, pp. 75–84.
- [8] Z. Feng, B. Xue, C. Wang, and F. Zhou, “Safe and socially compliant robot navigation in crowds with fast-moving pedestrians via deep reinforcement learning,” *Robotica*, pp. 1–19, 2024.
- [9] A. Francis, C. Pérez-d’Arpino, C. Li, F. Xia, A. Alahi, R. Alami, A. Bera, A. Biswas, J. Biswas, R. Chandra *et al.*, “Principles and guidelines for evaluating social robot navigation algorithms,” *arXiv preprint arXiv:2306.16740*, 2023.
- [10] M. Sebastian, S. B. Banisetty, and D. Feil-Seifer, “Socially-aware navigation planner using models of human-human interaction,” in *2017 26th IEEE international symposium on robot and human interactive communication (RO-MAN)*. IEEE, 2017, pp. 405–410.
- [11] Z. Henkel, C. L. Bethel, R. R. Murphy, and V. Srinivasan, “Evaluation of proxemic scaling functions for social robotics,” *IEEE Transactions on Human-Machine Systems*, vol. 44, no. 3, pp. 374–385, 2014.
- [12] S. Saunderson and G. Nejat, “How robots influence humans: A survey of nonverbal communication in social human–robot interaction,” *International Journal of Social Robotics*, vol. 11, no. 4, pp. 575–608, 2019.
- [13] S. Poddar, C. Mavrogiannis, and S. S. Srinivasa, “From crowd motion prediction to robot navigation in crowds,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 6765–6772.
- [14] M. Dharmadhikari, T. Dang, L. Solanka, J. Loje, H. Nguyen, N. Khedekar, and K. Alexis, “Motion primitives-based path planning for fast and agile exploration using aerial robots,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 179–185.
- [15] E. Leurent, “An environment for autonomous driving decision-making,” <https://github.com/eleurent/highway-env>, 2018.