

Effects of Human Motion Prediction Quality on Robot Navigation and Human Impressions in Teamwork Scenarios

Andrew Stratton¹, Phani Teja Singamaneni², Rachid Alami², and Christoforos Mavrogiannis¹

Abstract—Human motion prediction is critical for ensuring the seamless integration of mobile robots in dynamic human environments like warehouses, hospitals, and manufacturing plants. While state-of-the-art models are highly performant on tracking human behavior in offline datasets, their integration on mobile robots operating in complex environments gives rise to deployment challenges like accounting for out-of-distribution instances and human perceptions. In this work, we methodically investigate the implications of human motion prediction on robot navigation performance, human productivity, and human perceptions in a scenario involving navigation among two human subjects in a constrained workspace. Through a large-scale user study at two different sites (currently in progress), we expect to uncover practical insights and future directions for embodying human motion prediction models on mobile robots operating in dynamic environments.

I. INTRODUCTION

Autonomously navigating mobile robots have the potential to complete a wide variety of tasks in a wide variety of human environments. Algorithms for safely and efficiently navigating around humans, referred to as *Social Robot Navigation (SRN)* algorithms, can often be decomposed into a *prediction* task and a *control* task [16]. Existing methods for prediction include optimization-based models which leverage graphs [23, 24], physics-inspired forces [9], or game-theory [27] along with assumptions such as cooperative collision-avoidance behavior and rationality to predict optimal future trajectories for humans, and learning-based models which leverage datasets of humans navigating crowded spaces to learn to predict future trajectories directly from histories of states [6, 20, 21], using a learned reward function [31], or using learned parameters of a physics model [30]. The predictions are then integrated with a control method, most commonly model-predictive control (MPC) [15, 18] or reinforcement learning (RL) [2, 3, 12, 13].

Prediction model use in SRN allows robots to execute anticipatory behaviors, enabling them to navigate more safely and fluently around users. However different models make different sets of assumptions that affect their predictions, and by extension, the behavior of the robot. For instance, models based on cooperative collision avoidance assume that humans will accept part of the responsibility for conflict resolution [17], however when distracted or rushing this may not always be the case. In the case of learning-based predictors, the distribution of scenarios encountered at deployment is assumed to be sufficiently similar to those of the training data, however

existing datasets are typically collected in one or a few fixed areas which limit the diversity of interactions [6, 20, 21]. SRN algorithms are promised to be deployed in many substantively different scenarios which admit distinct navigation behaviors, thus it remains unclear whether models trained on existing datasets are sufficient for deployment across varied environments, or whether alternative solutions such as scenario-specific datasets must be employed [26].

In this work, we investigate the effect of various existing prediction strategies on users interacting with a socially navigating mobile robot. We instantiate a collaborative workspace, which mimics many downstream environments for mobile robot deployment, and are completing a user study in which we hypothesize that increasing prediction accuracy will lead to improved task performance, increased user comfort, and decreased physical and cognitive load on users.

While prior laboratory studies on social robot navigation and human motion prediction exist, this work sets itself apart by focusing on prediction rather than control strategies [14], and exploring several different models rather than focusing the investigation on a single algorithm [15, 18, 29]. Furthermore, we are conducting a multi-site investigation, mirroring our experiments at two sites (University of Michigan and LAAS-CNRS) with two substantially different robot embodiments (Hello Robot Stretch and Willow Garage PR2, see Fig. 1). This will allow us to investigate differences in experimental outcomes resulting from robot platform size.

The study is still in progress, however we anticipate the results will provide insight into the design of human trajectory prediction algorithms for use in mobile robots, as well as the development of control algorithms for mobile robots leveraging human trajectory prediction models.

II. PROBLEM SETTING AND ALGORITHMS

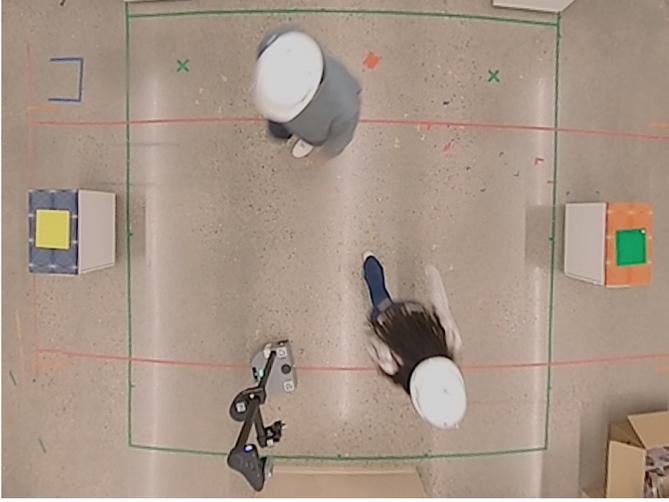
In order to study the use of distinct motion prediction models in social robot navigation, we desire a social navigation controller that allows for integrating several representative prediction models into it. To that end we instantiate a general-purpose social navigation model predictive controller, and define the integration of prediction models into its cost function. We then detail a set of prediction models operating at several levels of sophistication, which result in distinct social navigation behaviors without prediction-model specific tuning of the controller.

A. Problem Statement

We consider a robot navigating among $n \geq 1$ human agents in a workspace $\mathcal{W} \subseteq \mathbb{R}^2$ with a set of static obstacles $\mathcal{W}_{obs} \subseteq \mathcal{W}$. The robot starts from an initial configuration s_R and moves

¹Andrew Stratton and Christoforos Mavrogiannis are with the Department of Robotics, University of Michigan, Ann Arbor, USA. {arstr, cmavro}@umich.edu.

²Phani Teja Singamaneni and Rachid Alami are with LAAS-CNRS, Universite de Toulouse, CNRS, Toulouse, France. {ptsingaman, rachid.alami}@laas.fr.



(a) University of Michigan.



(b) LAAS-CNRS.

Fig. 1: Footage from our experiments at the University of Michigan (a) and LAAS-CNRS (b). In this work, we seek to investigate the implications of human motion prediction quality for robot navigation performance and human impressions during sustained close-quarters interaction.

towards a goal g_R by following a policy π_R , while humans navigate from their initial configurations s_i towards their goals, g_i by following a policy π_i , $i \in \mathcal{N}$; agents' goals are unknown to one another. The robot occupies an area $A_R \in \mathcal{W}$, and each human occupies an area $A_i \in \mathcal{W}$. We assume that the robot obtains observations of the world state, consisting of observations of human and static obstacle positions $(s_R^t, s_{1:n}^t, \mathcal{W}_{obs})$, and uses a length- h history of observations to determine controls $u_R^t = \pi_R(s_R^{t-h:t}, s_{1:n}^{t-h:t}, \mathcal{W}_{obs})$. Our goal is to control the robot such that it reaches goal g_R safely and efficiently, while also abiding by social norms.

B. Social Robot Navigation as Model Predictive Control

One way to determine π_R is to formulate social navigation as a discrete-time optimal control problem and solve for π_R^t using a model predictive control framework. We employ the following formalization:

$$\begin{aligned} \mathbf{u}^* &= \underset{\mathbf{u} \in \mathcal{U}}{\operatorname{argmin}} \mathcal{J}(\mathbf{s}_R, \mathbf{s}_{1:n}) \\ \text{s.t. } s_R^{t+1} &= g(s^t, u^t) \\ \mathbf{s}_{1:n} &= f(s_R^{t-h:t}, s_{1:n}^{t-h:t}, \mathcal{W}_{obs}), \end{aligned} \quad (1)$$

where \mathbf{s}_R is a robot state rollout obtained by evaluating prospective control sequence \mathbf{u} from control space \mathcal{U} using dynamics function g , $\mathbf{s}_{1:n}$ is a set of trajectory predictions for co-navigating humans obtained using prediction model f , and \mathcal{J} is a cost function capturing obstacle avoidance, social awareness, and goal-directedness. We use the vanilla-MPC cost from [15] augmented with an obstacle avoidance term:

$$\mathcal{J}(\mathbf{s}_R, \mathbf{s}_{1:n}, \mathcal{W}_{obs}) = a_g \mathcal{J}_g(\mathbf{s}_R) + a_d \mathcal{J}_d(\mathbf{s}, \mathbf{s}_{1:n}) + a_o \mathcal{J}_o(\mathbf{s}_R, \mathcal{W}_{obs}) \quad (2)$$

Where \mathcal{J}_g penalizes distance to the goal:

$$\mathcal{J}_g(\mathbf{s}_R) = \sum_{k=0}^{N-1} (s_R^{k+1} - g_R)^T Q_g (s_R^{k+1} - g_R) \quad (3)$$

\mathcal{J}_d penalizes distance to other agents via Kirby's *Asymmetric Gaussian Integral Function* [10]:

$$\mathcal{J}_d(\mathbf{s}_R, \mathbf{s}_{1:n}) = \sum_{k=0}^{N-1} \sum_{i=1}^n A_d^2(s_R^{k+1}, s_i^{k+1}) \quad (4)$$

\mathcal{J}_o penalizes penetration of obstacles:

$$\mathcal{J}_o(\mathbf{s}_R, \mathcal{W}_{obs}) = \sum_{k=0}^{N-1} \mathbb{1}(A_R(s_R^{k+1}) \cap \mathcal{W}_{obs} \neq \emptyset) \quad (5)$$

and a_g , a_d , and a_o are constants controlling the relative importance of each cost. Agent state transitions in (2) are approximated using a constant velocity model.

C. Human Trajectory Prediction

To obtain the human trajectory predictions $\mathbf{s}_{1:n}$ in (1), we use functions f which uses a length- h (note h can be 0) window of the history of environment observations $(s_R^{t-h:t}, s_{1:n}^{t-h:t}, \mathcal{W}_{obs})$ to predict finite horizon length- T joint-trajectory predictions. Specifically, we investigate the following five models:

No model: The robot does not perceive the humans navigating around it, that is $n = 0$, and thus $\mathcal{J}_d(\cdot, \cdot) = 0$.

Static: Humans are predicted to stay in place over the full time horizon, that is $f(s_R^{t-h:t}, s_{1:n}^{t-h:t}, \mathcal{W}_{obs})^{t'} = (s_R^t, s_{1:n}^t)$, $t' \in \{t+1, \dots, T\}$.

Constant Velocity: Humans are predicted to continue moving at their current velocity, that is

$$\begin{aligned} d_i &= (s_i^t - s_i^{t-1}) \\ f(s_R^{t-h:t}, s_{1:n}^{t-h:t}, \mathcal{W}_{obs})_{i}^{t'} &= s_i^t + d_i \cdot (t' - t), \\ t' &\in \{t+1, \dots, T\} \end{aligned} \quad (6)$$

Human Scene Transformer: The Human Scene Transformer (HST) [21] is a transformer-based generative model which outputs a Gaussian Mixture Model distribution for human i 's position at future timestep t' as follows:

$$P(s_i^{t'} | s_R^{t-h:t}, s_{1:n}^{t-h:t}) = \sum_{m=1}^M w_m \mathcal{N}(s; \sigma_{m,i}^{t'}, \mu_{m,i}^{t'}) \quad (7)$$

where m is the number of trajectory modes predicted, w_m is a predicted likelihood weight for each mode, and σ, μ are parameters predicted per-mode. To obtain the final prediction, we sample a single trajectory from the most likely mode.

CoHAN: CoHAN [23–25] is a cooperative human-aware navigation framework that uses a graph optimization to plan the robot and human trajectories together. Different prediction methodologies can be used in CoHAN, and for this study, we chose the constant velocity prediction modality. When CoHAN receives a goal for the robot, it extracts global paths for the robot (through A*) and the humans (by interpolating positions using constant velocity) in the given environment map. The optimization framework takes these global paths, map data and different social constraints into account while planning the trajectories for the robot and the humans. The planned trajectories for humans can be seen as predictions, and for a given human i , the trajectory, $\Gamma(s_i^t, t)$, can be obtained as:

$$\begin{aligned} d_i &= (s_i^t - s_i^{t-1}) \\ p_i^t &= s_i^t + d_i(t_k - t), \\ t_k &\in \{t + 1, \dots, T_n\} \\ \Gamma(s_i^t, t) &= g(p_i^t, s_i^t, g_R, map) \end{aligned} \quad (8)$$

where p_i is the global path, T_n is the prediction time inside CoHAN, and g represents the optimization result that takes robot’s goal, g_R and map data along with position data. Finally, the prediction for human i is obtained by sampling the trajectory, $\Gamma(s_i^t, t)$, at the right time using a function κ as follows:

$$f(s_R^{t-h:t}, s_{1:n}^{t-h:t}, \mathcal{W}_{obs})_i^{t'} = \kappa(\Gamma(s_i^t, t), t'), \quad (9)$$

$$t' \in \{t + 1, \dots, T\}$$

Although the workspace used in the study has a boundary, none of the existing prediction models are explicitly boundary aware. To prevent the predictions from crossing into obstacles, we truncate all predictions at the edges of the workspace by replacing all out-of-bounds predicted states with the latest state in the prediction still inside.

III. STUDY DESIGN

In order to evaluate the suite of algorithms from the previous section, we are running a within-subjects study (IRB HUM00259961) deploying the algorithms in a close-quarters collaborative assembly environment. Each experiment has two participants, who work together in a shared space to construct a set of towers, while the robot navigates the workspace alongside them. An identical setup is instantiated in two labs (University of Michigan and LAAS-CNRS) with the singular difference between the two being the robot platform used (the Hello Robot Stretch 2, and Willow Garage PR2, respectively). We completed 20 sessions at each site, resulting in a total of 80 participants, recruited from University students. We aim to compare the results in the two settings to investigate whether the robot embodiment affects the results.

Workspace Setup: We designed a workspace consisting of a $3.5 \times 3.5m$ square, with six workstations positioned at the perimeter of the space. Two stations at the bottom of the

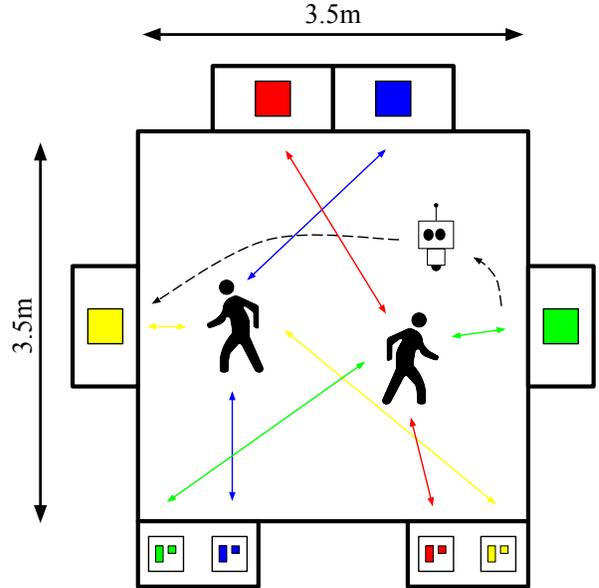


Fig. 2: Experimental setup. Each participant is responsible for two colors of towers (red and green or yellow and blue), while the robot traverses the workspace between stations. The ordering is fixed so the green and yellow towers are completed before the red and blue.

workspace each have two color-coded tower assembly areas (resulting in four distinct colors), each of which begins with one small and one large block already placed, on top of which participants complete two towers. Each of the remaining four stations have piles of ten bricks, five large and five small, corresponding to a single colored assembly station. The setup is visualized in Fig. 2.

Participant Tasks: Participants are tasked with constructing four towers each, corresponding to large and small towers for two colors. The participants are instructed to walk naturally, and complete each tower one block at a time, meaning each of them had to traverse the workspace 40 times to complete a single trial. To ensure participants maintained a consistent strategy regarding the order of tower completion, they were instructed to complete the towers in a fixed order (again visualized in Fig. 2). Participants begin at the bottom of the workspace by standing in front of the two assembly stations, and the task is considered complete when they have returned to their starting positions.

Robot Task: Over the course of a single trial, the robot navigates between the stations in a fixed order, running one of the tested algorithms. The order was randomly generated before experiments began, with the constraint that the robot’s previous and next goals are always on opposite sides of the workspace. The robot continues to navigate until both participants have returned to their starting positions, at which point the robot is stopped in place. Participants are specifically told that the robot is also part of their team, and that its task is to supervise progress of the assembly, with the specific goal of visiting as many workstations as possible.

Algorithm Implementation Details: We instantiate the control algorithm from section II using a Model Predictive Path Integral (MPPI) Controller [11, 28]. With a GPU par-

allelized implementation, the controller runs at 50 HZ with a history of $h = 8$, horizon of $T = 12$, Δt of 0.4s, and 500 rollouts. *HST* and *CoHAN* run at slower frequencies (20 HZ and 8 HZ respectively) and thus asynchronously update the trajectory predictions used in the controller. For the *HST*, we use a checkpoint trained on the ETH/UCY dataset which predicts $m = 20$ modes.

Metrics: Based on literature [4, 5, 16], we use the following performance metrics. We also note that *participants* refers to the recruited human subjects, and does not include the robot:

- *Participant Task Completion Time:* The average amount of time for each participant to complete their task, used to measure participant task completion efficiency.
- *Participant path irregularity [7]:* The amount of unnecessary turning per unit path length for participants, measured in $\frac{rad}{m}$, calculated as $\sum_{Path} \frac{Rotation - Min. rotation needed}{Path length}$. Used as a proxy for participant workload.
- *Participant acceleration:* The average acceleration for each participant. Used as a proxy for user comfort.

We also introduce a new metric, *Team average goals per second*, which captures the efficiency of the entire team, including the robot. The inclusion of the robot is particularly important because measuring productivity with respect to only the humans admits a trivial solution in which the robot simply moves to a corner of the workspace and does not actually move between goals, thus minimizing its intrusion but at the expense of completing its own task. An ideal algorithm will not be disruptive and will still complete its own task effectively.

Surveys: We use two surveys to obtain user impressions:

- *Robot Social Attributes Scale (RoSAS) [1]:* A scale containing 18 questions measuring impressions of *discomfort, competence, and warmth*. We specifically use the six questions measuring *discomfort*.
- *NASA Task-Load Index (NASA TLX) [8]:* A survey for assessing the physical and cognitive load on users. We use all questions on the original 21 point scale.

Hypotheses. We expect that as the accuracy of the motion predictions increase, the robot will navigate more effectively, thus improving task completion and user perceptions. Additionally, we expect that the more imposing frame of the PR2 will be less comfortable for participants. Based on these expectations, We formulate the following hypotheses:

H1: Using socially aware navigation algorithms results in improved working conditions and outcomes.

- *1a)* Socially aware algorithms have higher team average goals per second than non-social algorithms.
- *1b)* Socially aware algorithms are associated with lower perceived workload on users than non-social algorithms, as measured by the NASA TLX survey and users’ average path irregularities.
- *1c)* Socially aware algorithms are associated with increased user perceptions of comfort than non-social algorithms, as measured by the RoSAS survey and users’ average accelerations.

H2: Using prediction models with lower average displacement error (ADE) results in improved working conditions and

outcomes.

- *1a)* Decreasing prediction model ADE is associated with decreasing task completion times.
- *1b)* Decreasing prediction model ADE is associated with decreasing perceived workload on users, as measured by the NASA TLX survey and users’ average path irregularities.
- *1c)* Decreasing prediction model ADE is associated with increased user perceptions of comfort than non-social algorithms, as measured by the RoSAS survey and users’ average accelerations.

H3: A larger robot embodiment is associated with decreased perceptions of user comfort than a smaller embodiment, as measured by the RoSAS survey and users’ average accelerations.

IV. RESULTS AND DISCUSSION

The study is still in progress at both sites. We intend to include data analysis, including the support level for each hypothesis.

Thus far, members of the study team have qualitatively observed that, as anticipated, *No Model* causes uncomfortable interactions by blindly moving towards participants. The effects are most noticeable when they are at stations and are occupied with their task, and thus unable to perceive the robot moving into their space until it is very close. The *Static* model has generally caused overly-conservative behavior, due to the robot often incorrectly predicting that its path will be blocked. *CV*, *CoHAN*, and *HST* have produced more balanced results, with efficiency closer to *No Model*, while still taking actions that reduce disruption of the participants. Additionally, we observe differences in mean discomfort levels between sites, which indicates the platform difference may be causing differences in user experience. Ultimately, we aim to provide a finer-grained analysis of the predictions and navigation behaviors produced by each method, and relate them to our collected objective and subjective metrics.

While social robot navigation is often cited as a downstream use-case for human motion prediction models, its integration, and particularly the integration’s effects on humans the robot interacts with, are not yet well understood. This study will delve into the benefits and limitations of leveraging learning and model-based trajectory prediction methods in an *embodied* setting, where the trajectories are not passively observed, but instead are used actively for decision making by an agent in the scenario. Furthermore, the two-site experiment will allow us to analyze the effects of the embodiment of the agent itself. Finally, while recent work has begun to take an interest in constrained, task-driven settings [19, 22, 29], most existing large-scale datasets and methods focus on open spaces. This study will provide insight into whether these approaches can be performant in closer-quarters, contextualized settings.

ACKNOWLEDGEMENTS

The authors would like to thank Pranav Goyal for helpful discussions on study design, and Mo Xu for his contributions to the motion prediction system implementation.

REFERENCES

- [1] C. M. Carpinella, A. B. Wyman, M. A. Perez, and S. J. Stroessner. The robotic social attributes scale (RoSAS): Development and validation. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction, HRI '17*, page 254–262, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450343367.
- [2] C. Chen, Y. Liu, S. Kreiss, and A. Alahi. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 6015–6022, 2019.
- [3] C. Chen, S. Hu, P. Nikdel, G. Mori, and M. Savva. Relational graph learning for crowd navigation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10007–10013, 2020.
- [4] A. Francis et al. Principles and guidelines for evaluating social robot navigation algorithms. *arXiv:2306.16740 [cs.RO]*, 2023.
- [5] Y. Gao and C.-M. Huang. Evaluation of socially-aware robot navigation. *Frontiers in Robotics and AI*, 8, 01 2022.
- [6] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi. Social GAN: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2255–2264, 2018.
- [7] J. Guzzi, A. Giusti, L. M. Gambardella, and G. Di Caro. Human-friendly robot navigation in dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 423–430, 05 2013.
- [8] S. G. Hart and L. E. Staveland. Development of NASA-TLX (task load index): Results of empirical and theoretical research. *Human Mental Workload*, 1(3):139–183, 1988.
- [9] D. Helbing and P. Molnár. Social force model for pedestrian dynamics. *Physical Review E*, 51(5):4282–4286, 1995.
- [10] R. Kirby. *Social Robot Navigation*. PhD thesis, Carnegie Mellon University, 2010.
- [11] A. Lab. Pytorch mppi implementation, 2024.
- [12] S. Liu, P. Chang, W. Liang, N. Chakraborty, and K. Driggs-Campbell. Decentralized structural-rnn for robot crowd navigation with deep reinforcement learning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3517–3524, 2021.
- [13] S. Liu, P. Chang, Z. Huang, N. Chakraborty, K. Hong, W. Liang, D. L. McPherson, J. Geng, and K. Driggs-Campbell. Intention aware robot crowd navigation with attention-based interaction graph. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 12015–12021, 2023.
- [14] C. Mavrogiannis, A. M. Hutchinson, J. Macdonald, P. Alves-Oliveira, and R. A. Knepper. Effects of distinct robot navigation strategies on human behavior in a crowded environment. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 421–430, 2019.
- [15] C. Mavrogiannis, K. Balasubramanian, S. Poddar, A. Gandra, and S. S. Srinivasa. Winding through: Crowd navigation via topological invariance. *IEEE Robotics and Automation Letters*, 8(1):121–128, 2023.
- [16] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh. Core Challenges of Social Robot Navigation: A Survey. *Transactions on Human-Robot Interaction*, 12(3), 2023.
- [17] C. I. Mavrogiannis, W. B. Thomason, and R. A. Knepper. Social momentum: A framework for legible navigation in dynamic multi-agent environments. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 361–369, 2018.
- [18] S. Poddar, C. Mavrogiannis, and S. S. Srinivasa. From crowd motion prediction to robot navigation in crowds. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6765–6772, 2023.
- [19] A. Rudenko, T. P. Kucner, C. S. Swaminathan, R. T. Chadalavada, K. O. Arras, and A. J. Lilienthal. Thör: Human-robot navigation data collection and accurate motion trajectories dataset. *IEEE Robotics and Automation Letters*, 5(2):676–682, 2020.
- [20] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 683–700, 2020.
- [21] T. Salzmann, L. Chiang, M. Ryll, D. Sadigh, C. Parada, and A. Bewley. Robots that can see: Leveraging human pose for trajectory prediction. *IEEE Robotics and Automation Letters*, 8(11):7090–7097, 2023.
- [22] T. Schreiter, T. Rodrigues de Almeida, Y. Zhu, E. Gutierrez Maestro, L. Morillo-Mendez, A. Rudenko, L. Palmieri, T. P. Kucner, M. Magnusson, and A. J. Lilienthal. ThÖr-magni: A large-scale indoor motion capture recording of human movement and robot interaction. *The International Journal of Robotics Research*, page 02783649241274794.
- [23] P. T. Singamaneni and R. Alami. Hateb-2: Reactive planning and decision making in human-robot co-navigation. pages 179–186, 08 2020. doi: 10.1109/RO-MAN47096.2020.9223463.
- [24] P. T. Singamaneni, A. Favier, and R. Alami. Human-aware navigation planner for diverse human-robot interaction contexts. pages 5817–5824, 09 2021. doi: 10.1109/IROS51168.2021.9636613.
- [25] P. T. Singamaneni, A. Favier, and R. Alami. Watch out! there may be a human. addressing invisible humans in social navigation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11344–11351, 2022.
- [26] A. Stratton, K. Hauser, and C. Mavrogiannis. Characterizing the complexity of social robot navigation scenarios, 2024.

- [27] M. M. Sun, F. Baldini, K. Hughes, P. Trautman, and T. Murphey. Mixed strategy nash equilibrium for crowd navigation. *The International Journal of Robotics Research*, 2024. doi: 10.1177/02783649241302342.
- [28] G. Williams, A. Aldrich, and E. Theodorou. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*, 40:1–14, 01 2017. doi: 10.2514/1.G001921.
- [29] M. S. Yasar, M. M. Islam, and T. Iqbal. Posetron: Enabling close-proximity human-robot collaboration through multi-human motion prediction. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction, HRI '24*, page 830–839, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400703225. doi: 10.1145/3610977.3635006.
- [30] J. Yue, D. Manocha, and H. Wang. Human trajectory prediction via neural social physics. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022.
- [31] B. D. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. A. Bagnell, M. Hebert, A. K. Dey, and S. Srinivasa. Planning-based prediction for pedestrians. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3931–3936, 2009.